

# On Dynamic Flow-Sensitive Floating-Label Systems

Pablo Buiras<sup>1</sup>, Deian Stefan<sup>2</sup>, and Alejandro Russo<sup>1</sup>

<sup>1</sup> Chalmers University of Technology

<sup>2</sup> Stanford University

**Abstract.** Flow-sensitive analysis for information-flow control (IFC) allows data structures to have mutable security labels, i.e., labels that can change over the course of the computation. This feature is often used to boost the permissiveness of the IFC monitor, by rejecting fewer runs of programs, and to reduce the burden of explicit label annotations. However, adding flow-sensitive constructs (e.g., references or files) to a dynamic IFC system is subtle and may also introduce high-bandwidth covert channels. In this work, we extend LIO—a language-based floating-label system—with flow-sensitive references. The key insight to safely manipulating the label of a reference is to not only consider the label on the data stored in the reference, i.e., the reference label, but also the *label on the reference label* itself. Taking this into consideration, we provide an *upgrade* primitive that can be used to change the label of a reference in a safe manner. We additionally provide a mechanism for automatic upgrades to eliminate the burden of determining when a reference should be upgraded. This approach naturally extends to a concurrent setting, which has not been previously considered by dynamic flow-sensitive systems. For both our sequential and concurrent calculi we prove non-interference by embedding the flow-sensitive system into the original, flow-insensitive LIO calculus—a surprising result on its own.

## 1 Introduction

Modern software systems are composed of many complex components that handle sensitive data. In many cases (e.g., mobile and web applications) these disparate components are provided by different authors, of varying trustworthiness. Unfortunately, because today’s software development tools do not provide a means for protecting sensitive data from untrusted code, data theft and corruption is prevalent.

Information-flow control (IFC) is a promising approach to security that provides data confidentiality and integrity in the presence of untrusted code. At a high level, IFC tracks and controls the flow of information through a system according to a security policy, usually *non-interference* [15]. Non-interference states that public events should not depend on sensitive data and dually, trusted data should not be affected by untrusted events. Hence, with IFC, the program is guaranteed to preserve data confidentiality and integrity, even when composed

of untrusted components. Indeed, this appealing guarantee has recently led to significant research and development efforts that use IFC to secure web applications (e.g., [9, 14, 17, 41]) and mobile platforms (e.g., [11, 20]).

To ensure data confidentiality and integrity, these dynamic IFC systems associate *security labels* with data and monitor where such data can flow [25, 37]. In this paper, we use the labels  $\mathbf{H}$  and  $\mathbf{L}$ , to respectively denote secret and public data, and ensure that information cannot flow from a secret entity into a public one, i.e., the labels are ordered such that  $\mathbf{L} \sqsubseteq \mathbf{H}$  and  $\mathbf{H} \not\sqsubseteq \mathbf{L}$ . In general, the partial order  $\sqsubseteq$  (label check) is used to govern the allowed flows. We remark that our results apply to arbitrary lattices that may also express integrity concerns [25, 37], we only use the two-point lattice for simplicity of exposition.

One of the facets of IFC analysis lies in how such labels, when associated with objects, are treated [19]. Specifically, some IFC systems (e.g., [8, 10, 18, 21, 35, 36, 44]) treat labels on objects as *immutable* and do not allow for changes over the lifetime of the program, i.e., labels of objects are *flow-insensitive*. In contrast, other systems (e.g., [2, 3, 43]) are *flow-sensitive*, i.e., they allow object labels to change, in certain conditions, according to the sensitivity of the data that is stored in the object. In general, these flow-sensitive systems are more permissive, i.e., they allow programs that flow-insensitive monitors would reject.

Consider, for instance, a web application that writes to a labeled log while servicing user requests. If the label of the log is  $\mathbf{L}$ , a flow-insensitive IFC monitor would disallow writing any sensitive data (e.g., error messages containing user-supplied data) to the log, since this would constitute a leak. However, in a flow-sensitive system, the label of the log can change (to  $\mathbf{H}$ ), as to accommodate the kinds of data being written to the log. For many applications, allowing labels to change in such a way is very desirable—it alleviates the burden of having to a-priori determine the precise labels of objects (e.g., the log).

Unfortunately, naively introducing flow-sensitive objects to a dynamic IFC system can turn label changes into a covert channel [31]. Consider the code fragment of Figure 1 where references  $l$  and  $h$  are respectively labeled  $\mathbf{L}$  and  $\mathbf{H}$ . By naively allowing arbitrary label changes—even if the new label is more restricting—we can leak the contents of  $h$  into  $l$ . In particular, suppose that the temporary variable  $tmp$  is initially labeled  $\mathbf{L}$ . If the value stored in  $h$  is *True*, then in the first

```

 $l$       := True
 $tmp$  := False
if  $h$  then  $tmp$  := True
if  $\neg tmp$  then  $l$  := False

```

Fig. 1: Flow-sensitive attack

conditional we assign *True* into  $tmp$  and raise its label to  $\mathbf{H}$ , reflecting the fact that the branch condition depends on sensitive data. Since the  $tmp$  is *True*, the branch condition for the second conditional is *False* and thus the value and label of  $l$  are left intact, i.e., *True* at  $\mathbf{L}$ . However, if  $h$  is *False*, then the value and label of  $tmp$  do not change—the first assignment is not executed. Instead, the second assignment, which sets  $l$  to *False*, is performed; since the label of the branch condition is  $\mathbf{L}$ , the label of  $l$  remains  $\mathbf{L}$ . Note that in both cases the label of  $l$

stays **L**, but the value of  $l$  is the same as the secret  $h$ . (Hence why *label change* is considered a covert channel.) In systems such as LIO and Breeze, which allow labels to be inspected, this attack can be further simplified by simply checking the label of  $tmp$  after the first assignment—if the secret is true then the label will be **H**, otherwise it will be **L**.

This attack is not new, and, to ensure that the covert channel is not introduced when adding flow-sensitive references in such a way, several solutions have already been proposed. These solutions fall into roughly three categories. First, the IFC monitor can incorporate static information to ensure that such leaks are disallowed [31]. Second, the IFC monitor can forbid certain label changes, depending on the context (e.g., the program counter ( $pc$ ) label [33]). For instance, the *no-sensitive upgrades* policy disallows raising the label of a public reference in a sensitive context (e.g., when a branch condition is **H**) [2, 43]. And, third, the monitor can disallow branches that depend on certain variables, for which the label was mutated, as done by the *permissive upgrades* policy [3].

In this paper, we take a fresh perspective on flow-sensitivity in the context of coarse-grained floating-label systems, in particular, the LIO IFC system [35, 36]. LIO brings ideas from IFC Operating Systems—notably, HiStar [44]—into a language-based setting. In particular, LIO takes an OS-like coarse-grained approach by associating a single “current” floating-label with a computation (and everything in scope), instead of heterogeneously labeling every variable, as typically done by language-based systems (e.g., [27, 34]). This floating-label is raised (e.g., from **L** to **H**) to accommodate reading sensitive data and thus serves as a form of “taint” reflecting the sensitive of data in context, i.e., LIO is flow-sensitive in the current label. (This can be seen as raising the  $pc$  in more traditional language-based systems.) In turn, the LIO monitor uses the “current” label to restrict where the computation can write (e.g., once the current label is raised to **H**, it can no longer write to references labeled **L**). However, like other IFC systems, LIO is *flow-insensitive* in object labels.

This work extends the LIO IFC system, both the sequential and concurrent versions, to incorporate flow-sensitive references. A key insight of this work is to consider labels of references as being composed of two elements: the reference label describing the confidentiality (integrity) of the stored value, and another label, called *the label on the label*, which describes the confidentiality (integrity) of the reference label itself. Our monitor, then only forbids changing a label of a reference if *the label on the label is below the current floating-label*. Inspired by [17], we add a primitive for safely and explicitly *upgrading* labels. This boosts the permissiveness of LIO, and, for instance, allows programs, such as the logging web application described above, which would otherwise be rejected by the IFC monitor.

To reduce the burden of introducing upgrade annotations, our calculus provides a means for automatically upgrading references for which the computation is about to “lose” write access, i.e., before tainting the computation by raising the current label, we first upgrade all the references whose labels are below the (new) current label. While secure, this feature facilitates a form of *label creep*,

wherein all flow-sensitive references might end up with labels that are “too high.” To further address this, we propose a block-structured primitive which only upgrades the labels of declared flow-sensitive references, while disallowing writes to undeclared ones.

By taking a fresh perspective on flow-sensitivity, we also show that our flow-sensitive extension can be entirely encoded using existing flow-insensitive constructs—the key insight is to explicitly model flow-sensitive values as *nested flow-insensitive labeled references*. In the context of LIO, this encoding has the added benefit of allowing us to prove non-interference by simply invoking previous results. Equally important, the sequential semantics for LIO with flow-sensitive references directly extend to the concurrent setting.

The contributions of this paper are as follows:

- We extend LIO to incorporate flow-sensitive objects, with a focus on references. Specifically, we introduce two explicit primitives to safely raise (**upgrade**) or downgrade (**downgrade**) the security label of references. This extension not only increases LIO’s permissiveness, but also provides a means for safely combining flow-insensitive and flow-sensitive references.
- We present a uniform treatment for flow-insensitive and flow-sensitive references in both sequential and concurrent settings. To the best of our knowledge, we are the first to analyze the challenges of purely dynamic monitors with flow-sensitive references in the presence of concurrency.
- A non-interference proof for the different calculi that leverages the encoding of flow-sensitive references using flow-insensitive constructs.

The novel aspect of this article, with respect to its conference version [7], is the extension of our formal results to consider a **downgrade** primitive that further boosts permissiveness. Additionally, we compare our approach with *no-sensitive-upgrade* [43] and *permissive-upgrade* [2]—two known policies for label changes.

We remark that while our development focuses on LIO, we believe that our results generalize to other sequential and concurrent floating-label systems (e.g., [10, 18, 21, 44]).

The rest of the paper is organized as follows. Section 2 provides an introduction to LIO and its formalization. Section 3 presents our flow-sensitivity extensions and enforcement mechanism. Section 4 extends this approach to the concurrent setting. Section 5 presents the embedding of our enforcement using flow-insensitive constructs, from which our formal security guarantees follow. We discuss related work in Section 7 and conclude in Section 8.

## 2 Introduction to LIO

LIO is a language-level IFC system, implemented as a library in Haskell. The library provides a new *monad*, *LIO*, atop which programmers implement computations, which may use the LIO API to perform side effects (e.g., mutate a reference or write to a file).

$$\begin{array}{lcl}
\text{Values } v & ::= & \text{True} \mid \text{False} \mid () \mid \lambda x.t \mid \ell \mid LIO^{\text{TCB}} t \\
\text{Terms } t & ::= & v \mid x \mid t \ t \mid \mathbf{fix} \ t \mid \mathbf{if} \ t \ \mathbf{then} \ t \ \mathbf{else} \ t \\
& & \mid t \otimes t \mid \mathbf{return} \ t \mid t \gg t \mid \mathbf{getLabel} \\
\text{Types } \tau & ::= & \text{Bool} \mid () \mid \tau \rightarrow \tau \mid \ell \mid LIO \ \tau \\
\text{Ops}_\ell \quad \otimes & ::= & \sqcup \mid \sqcap \mid \sqsubseteq
\end{array}$$

Fig. 2: Syntactic categories for base  $\lambda_\ell^{\text{uo}}$ .

The *LIO* monad implements a purely dynamic execution monitor. Specifically, *LIO* encapsulates the state necessary to enforce IFC for the computation under evaluation. Part of this state is the current (floating) label. Intuitively, the current label serves a role similar to the program counter (*pc*) of more-traditional IFC systems (e.g., [34]): it is used to restrict the current computation from performing side-effects that may compromise the confidentiality or integrity of data (e.g., by restricting where the current computation may write).

To soundly reason about IFC, every piece of data *must* be labeled, including literals, terms, and labels themselves. However and different from most language-based systems (e.g., [18, 26, 34]) where every value is explicitly labeled, *LIO* takes a coarse-grained approach and uses the current label to protect all values in scope. As in IFC operating systems [10, 44], in *LIO*, the current label  $l_{\text{cur}}$  is the label on all the non-explicitly labeled values in the context of a computation.

To allow for computations on resistive data, *LIO* raises the current label to protect newly read data. That is, the current label is raised to “float” above the labels of all the objects read by the current computation. Raising the current label allows computations to flexibly read data, at the cost of being more limited in where they can subsequently write. Concretely, a computation with current label  $l_{\text{cur}}$  can read data labeled  $l_d$  by raising its current label to  $l'_{\text{cur}} = l_{\text{cur}} \sqcup l_d$ , but can thereafter only write to entities labeled  $l_e$  if  $l'_{\text{cur}} \sqsubseteq l_e$ . Hence, for example, a public *LIO* computation can read secret data by first raising  $l_{\text{cur}}$  from **L** to **H**. Importantly, however, the new current label prevents the computation from subsequently writing to public entities.

## 2.1 $\lambda_\ell^{\text{uo}}$ : A coarse-grained IFC calculus

We give the precise semantics for *LIO* by extending the simply-typed, call-by-name  $\lambda$ -calculus; we call this extended IFC calculus  $\lambda_\ell^{\text{uo}}$ . The formal syntax of the core  $\lambda_\ell^{\text{uo}}$  calculus, parametric in the label type  $\ell$ , is given in Fig. 2. Syntactic categories  $v$ ,  $t$ , and  $\tau$  represent values, terms, and types, respectively. Values include standard primitives (Booleans, unit, and  $\lambda$ -abstractions) and terminals corresponding to labels ( $\ell$ ) and monadic values ( $LIO^{\text{TCB}} t$ ).<sup>3</sup> We note that values of the form  $LIO^{\text{TCB}} t$  denote computations subject to security checks. (In fact,

<sup>3</sup> We restrict our formalization to computations implemented in the *LIO* monad and only consider Haskell features relevant to IFC, similar to the presentation of *LIO* in [38].

$E ::= E \ t \mid \mathbf{fix} \ E \mid \mathbf{if} \ E \ \mathbf{then} \ t \ \mathbf{else} \ t \mid E \otimes t \mid v \otimes E$   
 $E ::= [] \mid E \mid E \gg t$

$$\begin{array}{c}
\text{GETLABEL} \\
\hline
\Sigma = (l_{\text{cur}}, \dots) \\
\hline
\langle \Sigma \mid E \mid \mathbf{getLabel} \rangle \longrightarrow \langle \Sigma \mid E \mid \mathbf{return} \ l_{\text{cur}} \rangle
\end{array}$$

Fig. 4: Evaluation contexts and **getLabel** reduction rule.

security checks are only applied to such values.) Terms are composed of standard constructs (values, variables  $x$ , function application, the **fix** operator, and conditionals), terminals corresponding to label operations ( $t \otimes t$ , where  $\sqcup$  is the join,  $\sqcap$  is the meet, and  $\sqsubseteq$  is the partial-order on labels), standard monadic operators (**return**  $t$  and  $t \gg t$ ), and **getLabel**, a term for inspecting the current label, as further explained below. We do not consider terms annotated with  $\cdot^{\text{TCB}}$  as part of the surface syntax, i.e., such syntax nodes are not made available to programmers and are solely used internally in our semantic description. Types consist of Booleans, unit, function types, labels, and *LIO* computations; since the  $\lambda_{\ell}^{\text{uo}}$  type system is standard, we do not discuss it further.

We include monadic terms in our calculus since (in Haskell) monads dictate the evaluation order of a program and encapsulate all side-effects, including I/O [24, 40]; LIO leverages monads to precisely control what (side-effecting) operations the programmer is allowed to perform at any given time. In particular, an LIO program is simply a computation in the *LIO* monad, composed from simpler monadic terms using *return* and *bind*. Term **return**  $t$  produces a computation which simply returns the value denoted by  $t$ . Term  $\gg$ , called *bind*, is used to sequence LIO computations. Specifically, term  $t \gg (\lambda x. t')$  takes the result produced by term  $t$  and applies function  $\lambda x. t'$  to it. (This operator allows computation  $t'$  to depend on the value produced by  $t$ .) We sometimes use Haskell's **do**-notation to write such monadic computations. For example, the term  $t \gg \lambda x. \mathbf{return} \ (x + 1)$ , which simply adds 1 to the value produced by the term  $t$ , can be written using **do**-notation as shown in Figure 3.

$\mathbf{do} \ x \leftarrow t$   
 $\mathbf{return} \ (x + 1)$

Fig. 3: **do**-notation

A top-level  $\lambda_{\ell}^{\text{uo}}$  computation is a *configuration* of the form  $\langle \Sigma \mid t \rangle$ , where  $t$  is the monadic term and  $\Sigma$  is the state associated with the term. As in [35, 36], we take an imperative approach to modeling the LIO state as a separate component of the configuration (as opposed to being part of the term). We partially define the state of  $\lambda_{\ell}^{\text{uo}}$  to (at least) contain the current label  $l_{\text{cur}}$ , i.e.,  $\Sigma = (l_{\text{cur}}, \dots)$ ; here,  $\dots$  denotes other parts of the state not relevant at this point. Under this definition, a top-level well-typed  $\lambda_{\ell}^{\text{uo}}$  term has the form  $\Delta, \Gamma \vdash t : \text{LIO } \tau$ , where  $\Delta$  is the store typing, and  $\Gamma$  is the usual type environment.

$$\begin{aligned}
v &::= \dots \mid Lb^{\text{TCB}} l t \\
t &::= \dots \mid \mathbf{label} \ t \ t \mid \mathbf{unlabel} \ t \mid \mathbf{labelOf} \ t \\
\tau &::= \dots \mid Labeled \ \tau \\
E &::= \dots \mid \mathbf{label} \ E \ t \mid \mathbf{unlabel} \ E \mid \mathbf{labelOf} \ E
\end{aligned}$$

$$\begin{array}{c}
\text{LABEL} \\
\hline
\frac{\Sigma = (l_{\text{cur}}, \dots) \quad l_{\text{cur}} \sqsubseteq l}{\langle \Sigma \mid \mathbf{E} [\mathbf{label} \ l \ t] \rangle \longrightarrow \langle \Sigma \mid \mathbf{E} [\mathbf{return} \ (Lb^{\text{TCB}} l t)] \rangle} \\
\\
\text{UNLABEL} \\
\hline
\frac{\Sigma = (l_{\text{cur}}, \dots) \quad l'_{\text{cur}} = l_{\text{cur}} \sqcup l \quad \Sigma' = (l'_{\text{cur}}, \dots)}{\langle \Sigma \mid \mathbf{E} [\mathbf{unlabel} \ (Lb^{\text{TCB}} l t)] \rangle \longrightarrow \langle \Sigma' \mid \mathbf{E} [\mathbf{return} \ t] \rangle} \\
\\
\text{LABELOF} \\
\hline
\frac{}{\mathbf{E} [\mathbf{labelOf} \ (Lb^{\text{TCB}} l t)] \longrightarrow \mathbf{E} [l]}
\end{array}$$

Fig. 5: Extending base  $\lambda_\ell^{\text{uo}}$  with labeled values.

We use evaluation contexts in the style of Felleisen and Hieb to specify the reduction rules for  $\lambda_\ell^{\text{uo}}$  [12]. Figure 4 defines the evaluation contexts for pure terms ( $E$ ) and monadic terms ( $\mathbf{E}$ ) for the base  $\lambda_\ell^{\text{uo}}$ . The definitions are standard; we solely highlight that monadic terms are evaluated only at the outermost use of bind ( $\mathbf{E} \gg t$ ), as in Haskell. For the base  $\lambda_\ell^{\text{uo}}$ , we also give the reduction rule for the monadic term **getLabel**, which simply retrieves the current label. As shown later, it is precisely this label that is used to restrict the reads/writes performed by the current computation. The rest of the reduction rules for the base calculus are straight forward and given Appendix A.

## 2.2 Labeled values

Using  $l_{\text{cur}}$  as the label on all terms in scope makes it trivial to deal with implicit flows. Branch conditions, which are simply values of type *Bool*, are already implicitly labeled with  $l_{\text{cur}}$ . Consequently, all the subsequent writes cannot leak this bit—the current label restricts all the possible writes. However, this coarse-grained labeling approach suffers from a severe restriction: a piece of code cannot, for example, write a public value (e.g., 42) to a public channel labeled **L** after observing secret data, even if the value is independent from the secret—once secret data is read, the current label is raised to **H** thereby “over tainting” the public data in scope.

To address this limitation, LIO provides *Labeled* values. A *Labeled* value is a term that is explicitly protected by a label, other than the current label. Figure 5 shows the extension of the base  $\lambda_\ell^{\text{uo}}$  with *Labeled* values.

The **label** terminal is used to explicitly label a term. As rule (LABEL) shows, **label**  $l \ t$  associates the supplied label  $l$  with term  $t$  by wrapping the term with

the  $Lb^{\text{TCB}}$  constructor. Importantly, it first asserts that the new label ( $l$ ), which will be used to protect  $t$ , is at least as restricting as the current label, i.e.,  $l_{\text{cur}} \sqsubseteq l$ .

Dually, terminal **unlabel** unwraps explicitly labeled values. As defined in rule (UNLABEL), given a labeled value  $Lb^{\text{TCB}} l t$ , **unlabel** returns the wrapped term  $t$ . Since the returned term is no longer explicitly labeled by  $l$ , and is instead protected by the current label,  $l_{\text{cur}}$  must be at least as restricting as  $l$ . To ensure this, the current label is raised from  $l_{\text{cur}}$  to  $l_{\text{cur}} \sqcup l$ , capturing the fact that the remaining computation might depend on  $t$ . This rule highlights the fact that the current label always “floats” above the labels of the values observed by the current computation.

The **labelOf** function provides a means for inspecting the label of a labeled value. As detailed by reduction rule (LABELOF), given a labeled value  $Lb^{\text{TCB}} l t$ , the function returns the label  $l$  protecting term  $t$ . This allows code to check the label of a labeled value before deciding to unlabel it, and thereby raising the current label. It is worth noting that regardless of the current label in the configuration, the label of a labeled value can be inspected—hence labels are effectively “public.”<sup>4</sup>

A common problem with dynamic IFC systems is *label creep* [33]—the raising of the current label to a point where the computation can no longer do anything useful. To avoid label creep, LIO provides **toLabeled** as a way to allow the current label to be *temporarily* raised during the execution of a given computation. We extend the terms and the pure evaluation context as  $t ::= \dots \mid \mathbf{toLabeled} \ t \ t$  and  $E ::= \dots \mid \mathbf{toLabeled} \ E \ t$ , respectively, and give the precise semantics for **toLabeled** as follows:

$$\text{TOLABELLED} \quad \frac{\langle \Sigma | t \rangle \longrightarrow^* \langle \Sigma' | LIO^{\text{TCB}} t' \rangle \quad \begin{array}{l} \Sigma = (l_{\text{cur}}, \dots) \quad l_{\text{cur}} \sqsubseteq l \\ \Sigma' = (l'_{\text{cur}}, \dots) \quad l'_{\text{cur}} \sqsubseteq l \end{array} \quad \Sigma'' = \Sigma \ltimes \Sigma'}{\langle \Sigma | \mathbf{E} [\mathbf{toLabeled} \ l \ t] \rangle \longrightarrow \langle \Sigma'' | \mathbf{E} [\mathbf{label} \ l \ t'] \rangle}$$

If the current label at the point of executing **toLabeled**  $l \ t$  is  $l_{\text{cur}}$ , **toLabeled** evaluates  $t$  to completion ( $\langle \Sigma | t \rangle \longrightarrow^* \langle \Sigma' | LIO^{\text{TCB}} t' \rangle$ ) and restores the current label to  $l_{\text{cur}}$ , i.e., **toLabeled** provides a separate context in which  $t$  is evaluated. (Here, the state merge function  $\ltimes$  is defined as:  $\Sigma \ltimes \Sigma' \triangleq \Sigma$ , in the next section we present an alternative definition.) We note that returning the result of evaluating  $t$  directly (e.g., as  $\langle \Sigma | \mathbf{E} [\mathbf{toLabeled} \ l \ t] \rangle \longrightarrow \langle \Sigma'' | \mathbf{E} [t'] \rangle$ ) would allow for trivial leaks; thus, **toLabeled** labels  $t'$  with  $l$  ( $\langle \Sigma'' | \mathbf{E} [\mathbf{label} \ l \ t'] \rangle$ ). This effectively states that the result of  $t$  is protected by label  $l$ , as opposed to the current label ( $l'_{\text{cur}}$ ) at the point  $t$  completed. Importantly, this requires that the result not be more sensitive than  $l$ , i.e.,  $l'_{\text{cur}} \sqsubseteq l$ .

<sup>4</sup> Since labeled values can be nested, this only applies to the labels of top-level labeled values. Indeed, even these labels are not public—they are protected by the current label. However, since code can always observe objects labeled at the current label, this is akin to being public.



$$\begin{array}{l}
v ::= \dots \mid Ref_{FI}^{TCB} \ l \ a \\
t ::= \dots \mid \mathbf{newRef}_s \ t \ t \mid \mathbf{writeRef}_s \ t \ t \mid \mathbf{readRef}_s \ t \\
\quad \mid \mathbf{labelOf}_s \ t \mid \mathbf{copyRef} \ t \ t \\
\tau ::= \dots \mid Ref_s \ \tau \\
E ::= \dots \mid \mathbf{newRef}_s \ E \ t \mid \mathbf{writeRef}_s \ E \ t \mid \mathbf{readRef}_s \ E \\
\quad \mid \mathbf{labelOf}_s \ E \mid \mathbf{copyRef} \ E \ t \mid \mathbf{copyRef} \ v \ E
\end{array}$$
  

$$\begin{array}{c}
\text{NEWREF-FI} \\
\hline
\Sigma = (l_{\text{cur}}, \mu_{FI}, \dots) \quad l_{\text{cur}} \sqsubseteq l \quad \mu'_{FI} = \mu_{FI} [a \mapsto Lb^{TCB} \ l \ t] \quad \Sigma' = (l_{\text{cur}}, \mu'_{FI}, \dots) \quad \text{fresh}(a) \\
\hline
\langle \Sigma \mid \mathbf{E} [\mathbf{newRef}_{FI} \ l \ t] \rangle \longrightarrow \langle \Sigma' \mid \mathbf{E} [\mathbf{return} (Ref_{FI}^{TCB} \ l \ a)] \rangle
\end{array}$$
  

$$\begin{array}{c}
\text{READREF-FI} \\
\hline
\Sigma = (l_{\text{cur}}, \mu_{FI}, \dots) \\
\hline
\langle \Sigma \mid \mathbf{E} [\mathbf{readRef}_{FI} (Ref_{FI}^{TCB} \ l \ a)] \rangle \longrightarrow \langle \Sigma \mid \mathbf{E} [\mathbf{unlabel} \ \mu_{FI} (a)] \rangle
\end{array}$$
  

$$\begin{array}{c}
\text{WRITEREF-FI} \\
\hline
\Sigma = (l_{\text{cur}}, \mu_{FI}, \dots) \quad l_{\text{cur}} \sqsubseteq l \quad \mu'_{FI} = \mu_{FI} [a \mapsto Lb^{TCB} \ l \ t] \quad \Sigma' = (l_{\text{cur}}, \mu'_{FI}, \dots) \\
\hline
\langle \Sigma \mid \mathbf{E} [\mathbf{writeRef}_{FI} (Ref_{FI}^{TCB} \ l \ a) \ t] \rangle \longrightarrow \langle \Sigma' \mid \mathbf{E} [\mathbf{return} ()] \rangle
\end{array}$$
  

$$\begin{array}{c}
\text{LABELOF-FI} \\
\hline
E [\mathbf{labelOf}_{FI} (Ref_{FI}^{TCB} \ l \ a)] \longrightarrow E [l]
\end{array}$$
  

$$\begin{array}{c}
\text{COPYREF} \\
\hline
\Sigma = (l_{\text{cur}}, \mu_{FI}, \dots) \quad l_1 \sqsubseteq l_2 \quad l_{\text{cur}} \sqsubseteq l_2 \\
Lb^{TCB} \ l_1 \ v_1 = \mu_{FI} (a_1) \quad \mu'_{FI} = \mu_{FI} [a_2 \mapsto Lb^{TCB} \ l_2 \ v_1] \quad \Sigma' = (l_{\text{cur}}, \mu'_{FI}, \dots) \\
\hline
\langle \Sigma \mid \mathbf{E} [\mathbf{copyRef} (Ref_{FI}^{TCB} \ l_1 \ a_1) (Ref_{FI}^{TCB} \ l_2 \ a_2)] \rangle \longrightarrow \langle \Sigma' \mid \mathbf{E} [\mathbf{return} ()] \rangle
\end{array}$$

Fig. 6: Extending  $\lambda_\ell^{\text{uo}}$  with references.

### 2.3 Labeled references

To complete the description of LIO, we extend the  $\lambda_\ell^{\text{uo}}$  calculus with mutable, flow-insensitive references. Conceptually, flow-insensitive references are simply mutable *Labeled* values. Like labeled values, the label of a reference is immutable and serves to protect the underlying term. The immutable label makes the semantics straightforward: writing a term to a reference amounts to ensuring that the reference label is as restrictive as the current label, i.e., the reference label must be above the current label; reading from a reference taints the current label with the reference label.

The syntactic extensions to our calculus are shown in Figure 6. We use meta-variable  $s$  to distinguish flow-insensitive (FI) and flow-sensitive (FS) productions—the latter are described in Section 3. We also extend configurations to contain a

reference (memory) store  $\mu_{\text{FI}}: \Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \dots)$ ;  $\mu_{\text{FI}}$  maps memory addresses—spanned over by metavariable  $a$ —to *Labeled* values.

When creating a flow-insensitive reference, **newRef<sub>FI</sub>**  $l\ t$  creates a labeled value that guards  $t$  with label  $l$  ( $Lb^{\text{TCB}}\ l\ t$ ) and stores it in the memory store at a fresh address  $a$  ( $\mu_{\text{FI}}[a \mapsto Lb^{\text{TCB}}\ l\ t]$ ). Subsequently, the function returns a value of the form  $Ref_{\text{FI}}^{\text{TCB}}\ l\ a$  which simply encapsulates the reference label and address where the term is stored. We remark that since any references created within a **toLabeled** block may outlive the **toLabeled** block computation, the merge function used in rule (TOLABELED) must also account for this, i.e.,  $(l_{\text{cur}}, \mu_{\text{FI}}, \dots) \times (l'_{\text{cur}}, \mu'_{\text{FI}}, \dots) = (l_{\text{cur}}, \mu'_{\text{FI}}, \dots)$ .

Rule (READREF-FI) gives the semantics for reading a labeled reference; reading the term stored at address  $a$  simply amounts to unlabeled the value  $\mu(a)$  stored at the underlying address (**unlabel**  $\mu_{\text{FI}}(a)$ ).

Terminal **writeRef<sub>FI</sub>** is used to update the memory store with a new term. Note that **writeRef<sub>FI</sub>** *leaves the label of the reference intact*, i.e., the label of a flow-insensitive reference is never changed, but, as rule (WRITEREF-FI) shows in turn, requires the current label to be below the reference label when performing the write ( $l_{\text{cur}} \sqsubseteq l$ ).

Terminal **labelOf<sub>FI</sub>** has the benefit of allowing code to always inspect the label of a reference.

Terminal **copyRef** is to copy the contents of one reference to another, without inspecting the contents of either reference. As given by rule (COPYREF), the function copies the contents of a labeled reference into another one, as long as the source-reference label ( $l_1$ ) flows to the target-reference label ( $l_2$ ) and the usual condition for writing to an entity with label  $l_2$  also holds ( $l_{\text{cur}} \sqsubseteq l_2$ ). Since the computation does not read the source reference, the current label remains unchanged. We remarks that while **copyRef** can be encoded using **toLabeled**, we introduce **copyRef** explicitly since the use of **toLabeled** is prohibited in concurrent settings and our results rely on such a feature in both contexts (see Section 5),

### 3 Flow-sensitivity extensions

The flow-insensitive references described in the previous section are inflexible. Consider, for example, an application that uses a reference as a log. Since the log may contain sensitive information, it is important that the reference be labeled. Equally important is to be able to read the log at any point in the program to, for instance, save it to a file. Although labeling the reference with the top element in the security lattice ( $\top$ ) would always allow writes to the log, and **toLabeled** can be used to read the log and then write it to a file, this is unsatisfactory: it assumes the existence of a top element, which in some practical IFC systems, including HiStar [44] and Hails [14], does not exist. Moreover, it almost always over-approximates the sensitivity of the log. Hence, for example, a computation that never reads sensitive data, yet wishes to read the log content as to send error message to a user over the network (e.g., as done in a web application) cannot

```

leakRef :: RefFS Bool → LIO Bool
leakRef href = do
  tmp ← newRef L ()
  toLabeled H $ do h ← readRef href
                  when h $ writeRef tmp ()
  return $ labelOf tmp ≡ H

```

Fig. 7: Attack on LIO with naive treatment of flow-sensitive references. We omit subscripts for clarity.

do so—LIO prevents the computation from reading the log, which results in the computation getting tainted by  $\top$ , and subsequently writing to the network.<sup>5</sup> It is clear that even for such a simple use case, having references with labels that vary according to the sensitivity of what is stored in the reference is useful.

However, naively implementing flow-sensitive references can effectively introduce label changes as a covert channel. Suppose that we allow for the label of a reference to be raised to the current label at the time of the **writeRef**. So, for example, if the label of our log reference is **L** and the computation has read sensitive data (such that the current label is **H**), subsequently writing to the log will raise the label of the reference to **H**. Unfortunately, while this may appear safe, as previously shown in [2, 3, 31], the approach is unsound.

The code fragment in Figure 7 defines a function, *leakRef*, that can be used to leak the contents of a reference by leveraging the newly introduced covert channel: the label of references. (In this and future examples we use function **when** to denote an **if** statement without the **else** branch and (\$) as lightweight notation for function application, i.e.,  $f \$ x$  is the same as  $f(x)$ .) To illustrate an attack, suppose that the current label is public (**L**) and *leakRef* is called with a secret (**H**) reference (*href*). *leakRef* first creates a public reference *tmp* and, then, within the **toLabeled** block—which is used to ensure that the current label remains **L**—the label of this reference is changed to **H** if the secret stored in *href* is *True*, and left intact (**L**) if the secret is *False*. The value stored in *href* is revealed by simply inspecting the label of the *tmp* reference.<sup>6</sup>

Fundamentally, the label protecting the *label* of an object, such as a reference or labeled value, is the current label  $l_{\text{cur}}$  at the time of creation. Hence, to modify the label of the object within some context (e.g., **toLabeled** block) wherein the current label is  $l'_{\text{cur}}$ , it must be the case that  $l'_{\text{cur}} \sqsubseteq l_{\text{cur}}$ , i.e., we must be able to write data at sensitivity level  $l'_{\text{cur}}$  into an entity—the label of the object—labeled  $l_{\text{cur}}$ . This restriction is especially important if  $l_{\text{cur}} \sqsubset l'_{\text{cur}}$  and we can restore the current label from  $l'_{\text{cur}}$  to  $l_{\text{cur}}$ , since a leak would then be observable within the program itself. In the case where the label of the object is immutable, as is the case for flow-insensitive references (and labeled values), this is not a concern:

<sup>5</sup> Here, as in most IFC systems, we assume the network is public.

<sup>6</sup> The use of **labelOf** is not fundamental to this attack and in Appendix B we show an alternative attack that does not rely on such label inspection.

even if the current label is raised to  $l'_{\text{cur}}$  and then restored to  $l_{\text{cur}}$ , we do not learn any information more sensitive than  $l_{\text{cur}}$ —the label of the label at the time of creation—by inspecting the label of the reference (or value): the label has not changed!

Thus, to extend LIO with flow-sensitive references, we must account for the label on the label of the reference at the time of creation,  $l_{\text{cur}}$ . (This label is, however, immutable.) In turn, when changing the label of the reference, we must ensure that no data from the context at the time of the change, whose label is  $l'_{\text{cur}}$ , is leaked into the label of the reference by ensuring that  $l'_{\text{cur}} \sqsubseteq l_{\text{cur}}$ , i.e., we can write data labeled  $l'_{\text{cur}}$  into the label that is labeled  $l_{\text{cur}}$ .

Formally, we extend the  $\lambda_{\ell}^{\text{LIO}}$  syntax and reduction rules as shown in Figure 8; we call this calculus  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$ . To create a flow-sensitive reference **newRef**<sub>FS</sub>  $l$   $t$  creates a labeled value that guards  $t$  with label  $l$  ( $Lb^{\text{TCB}} l$   $t$ ). Since we wish to allow programmers to modify the label  $l$  of the reference, we additionally store the label on  $l$ , i.e., the current label  $l_{\text{cur}}$ , by simply labeling the already-guarded term ( $\mu'_{\text{FS}} = \mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} l_{\text{cur}} (Lb^{\text{TCB}} l$   $t)]$ ), as shown in rule (NEWREF-FS). Primitive **newRef**<sub>FS</sub> returns a  $\text{Ref}_{\text{FS}}^{\text{TCB}} a$  which simply encapsulates the fresh reference address where the doubly-labeled term is stored. Different from the constructor  $\text{Ref}_{\text{FI}}^{\text{TCB}}$ , the constructor  $\text{Ref}_{\text{FS}}^{\text{TCB}}$  does not encapsulate the label of the reference. This is precisely because the label of a flow-sensitive reference is mutable and must be looked up in the store. As given by rule (LABELOF-FS), **labelOf**<sub>FS</sub> returns the label of the reference after raising the current label (with **unlabel**) to account for the fact that the label of the reference  $l'$  is a value at sensitivity level  $l$ , i.e., we raise the current label to the join of the current label and the label on the label.

The rule for reading flow-sensitive references is standard. As given by rule (READREF-FS), **readRef**<sub>FS</sub> simply raises the current label to the join of the reference label and label on the reference label ( $l \sqcup l'$ ) and returns the protected value. This reflects the fact that the computation is observing both data at level  $l$  (the label on the reference) and  $l'$  (the actual term).

The rule for writing flow-sensitive references deserves more attention. First, **writeRef**<sub>FS</sub> as given by rule (WRITEREF-FS), ensures that the current computation can write to the reference by checking that  $l_{\text{cur}} \sqsubseteq (l \sqcup l')$ . We impose this condition instead of the two conditions  $l_{\text{cur}} \sqsubseteq l$  and  $l_{\text{cur}} \sqsubseteq l'$ —which respectively check that the current computation can modify both, the label of the reference, and the reference itself—since it is more permissive, yet still safe. When imposing the two conditions independently, certain programs, such as the one given in Figure 9, would fail. In this program, we first create a flow-sensitive reference labeled **H** when the current label is **L** (and thus the label on **H** is **L**). Then, we raise the label by reading from the reference. Finally, we attempt to write to the reference.

```

do  $r \leftarrow \text{newRef}_{\text{FS}} \mathbf{H} ()$ 
  readRefFS  $r$ 
  writeRefFS  $r ()$ 

```

Fig. 9: Permissiveness test.

$$\begin{array}{lcl}
v ::= \dots & | & Ref_{\text{FS}}^{\text{TCB}} t \\
t ::= \dots & | & \text{upgrade}_{\text{FS}} t t \mid \uparrow \mid \text{downgrade}_{\text{FS}} t t \\
E ::= \dots & | & \text{upgrade}_{\text{FS}} E t \mid \text{upgrade}_{\text{FS}} v E \\
& & \text{downgrade}_{\text{FS}} E t \mid \text{downgrade}_{\text{FS}} v E
\end{array}$$

NEWREF-FS

$$\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad l_{\text{cur}} \sqsubseteq l \quad \text{fresh}(a) \quad \mu'_{\text{FS}} = \mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} l_{\text{cur}} (Lb^{\text{TCB}} l t)] \quad \Sigma' = (l_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}})}{\langle \Sigma | \mathbf{E} [\text{newRef}_{\text{FS}} l t] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\text{return} (Ref_{\text{FS}}^{\text{TCB}} a)] \rangle}$$

READREF-FS

$$\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l' t) \quad l'' = l \sqcup l'}{\langle \Sigma | \mathbf{E} [\text{readRef}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a)] \rangle \longrightarrow \langle \Sigma | \mathbf{E} [\text{unlabel} (Lb^{\text{TCB}} l'' t)] \rangle}$$

WRITEREF-FS

$$\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l' t') \quad l_{\text{cur}} \sqsubseteq (l \sqcup l') \quad \mu'_{\text{FS}} = \mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} l (Lb^{\text{TCB}} l' t)] \quad \Sigma' = (l_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}})}{\langle \Sigma | \mathbf{E} [\text{writeRef}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) t] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\text{return} ()] \rangle}$$

WRITEREF-FS-FAIL

$$\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l' t') \quad l_{\text{cur}} \not\sqsubseteq (l \sqcup l')}{\langle \Sigma | \mathbf{E} [\text{writeRef}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) t] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\text{unlabel} (Lb^{\text{TCB}} l \uparrow)] \rangle}$$

LABELOF-FS

$$\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l' t')}{\langle \Sigma | \mathbf{E} [\text{labelOf}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a)] \rangle \longrightarrow \langle \Sigma | \mathbf{E} [\text{unlabel} (Lb^{\text{TCB}} l l')] \rangle}$$

UPGRADEREF

$$\frac{l_{\text{cur}} \sqsubseteq l \quad \Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l'' v) \quad \mu'_{\text{FS}} = \mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} l (Lb^{\text{TCB}} (l'' \sqcup l') v)] \quad \Sigma' = (l_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}})}{\langle \Sigma | \mathbf{E} [\text{upgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) l'] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\text{return} ()] \rangle}$$

DOWNGRADEREF

$$\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l'' v) \quad l_{\text{cur}} \sqsubseteq l \quad \mu'_{\text{FS}} = \mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} l (Lb^{\text{TCB}} (l \sqcup (l'' \sqcap l')) \uparrow)] \quad \Sigma' = (l_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}})}{\langle \Sigma | \mathbf{E} [\text{downgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) l'] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\text{return} ()] \rangle}$$

Fig. 8:  $\lambda_{\ell, \text{FS}}^{\text{uo}}$ : extension of  $\lambda_{\ell}^{\text{uo}}$  with flow-sensitive references.

Under our semantics, this program behaves as expected; however, when imposing the two conditions independently, the write fails—the current label does not flow to the label on the label of the reference.

Another case for **writeRef**<sub>FS</sub> which we must handle is when current label does not flow to the join of the reference label, i.e.,  $l_{\text{cur}} \not\sqsubseteq l \sqcup l'$ , and the write is disallowed. If the semantics simply got stuck, the current label (at the point of the stuck term) would not reflect the fact that the success of applying such rule depends on the label  $l'$ , which is itself protected by  $l$ . Indeed, this might lead to information leaks and we thus provide an explicit rule, (WRITEREF-FS-FAIL), for this failure case that first raises the current label (via **unlabel**) to  $l$  and then diverges; in the rule,  $\uparrow$  represents a divergent term for which we do not provide a reduction rule.

Note that **writeRef**<sub>FS</sub> does not modify the label of the reference. This is, in part, because we wish to keep the difference between flow-insensitive and flow-sensitive references as small as possible. Instead, we provide **upgrade**<sub>FS</sub> precisely for this purpose; this primitive is used to raise the label of a reference. Rule (UPGRADEREF) is straight forward—it simply ensures that the current computation can modify the label of the reference by checking that the current label flows to the label on the label ( $l_{\text{cur}} \sqsubseteq l$ ). Similarly, **downgrade**<sub>FS</sub> is used to lower the label of the reference, destroying its contents, i.e., replacing its value with  $\uparrow$ . Rules (UPGRADEREF) and (DOWNGRADEREF) are analogous; the main difference is that the former uses the join operation to combine the old and new labels ( $l'' \sqcup l'$ ), whereas the latter uses the meet operation ( $l'' \sqcap l'$ ). The **downgrade**<sub>FS</sub> primitive is useful when one wishes to store information that is less sensitive into a reference. Both **upgrade**<sub>FS</sub> and **downgrade**<sub>FS</sub> highlight that it is safe to raise or lower the label of a flow-sensitive reference, if that the label on the label still flows to the final label in the nested  $Lb^{\text{TCB}}$  structure.

### 3.1 Automatic upgrades

We can use  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  to implement various applications that rely on flow-sensitive references, even those that rely on policies such as the popular no-sensitive upgrades [2]. Using  $\lambda_{\ell, \text{FS}}^{\text{uo}}$ , we can also safely implement our logging application using a flow-sensitive reference. Unfortunately, our system (and others like it) requires that we insert **upgrades** before we raise the current label so that it is possible to write references in a more-sensitive context, e.g., to modify a public reference after reading a secret. In the case of the logging example, we would need to upgrade the label before reading any sensitive data, if we later wish to write to the log.

We provide an extension to  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  that can be used to automatically upgrade references. This extension, called  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$ , is given in Figure 10. Intuitively, whenever the current label is about to be raised, we first upgrade all the references in the  $\mu_{\text{FS}}$  store and then raise the current label. Rule (UPGRADESTORE) upgrades every reference in the flow-sensitive store  $\mu_{\text{FS}}$  by executing  $t_1 \gg t_2 \gg \dots \gg t_n$ , where  $t_i = \mathbf{upgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a_i) l$ . Term  $t \gg t'$  is similar to bind except that it discards the result produced by  $t$ . Since **unlabel** is the only function that

$$\begin{array}{c}
\text{UPGRADESTORE} \\
\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu_{\text{FS}} = \{a_1 \mapsto v_1, \dots, a_n \mapsto v_n\} \quad t_i = \text{upgrade}_{\text{FS}}(Ref_{\text{FS}}^{\text{TCB}} a_i) \ l, i = 1, \dots, n}{\langle \Sigma | \mathbf{E} [\text{upgradeStore}_{\text{FS}} \ l] \rangle \longrightarrow \langle \Sigma | \mathbf{E} [t_1 \gg \dots \gg t_n] \rangle} \\
\\
\text{UNLABEL-AU} \\
\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad l'_{\text{cur}} = l_{\text{cur}} \sqcup l \quad \langle \Sigma | \text{upgradeStore}_{\text{FS}} \ l'_{\text{cur}} \rangle \longrightarrow^* \langle l_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}} | LIO^{\text{TCB}} () \rangle \quad \Sigma' = (l'_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}})}{\langle \Sigma | \mathbf{E} [\text{unlabel} (Lb^{\text{TCB}} \ l \ t)] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\text{return} \ t] \rangle}
\end{array}$$

Fig. 10:  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$ : Extending  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$  with auto-upgrades.

raises the current label, we augment the (UNLABEL) rule with (UNLABEL-AU), given in Figure 10. This ensures that as the computation progresses it does not “lose” write access to its references. Returning to our logging example, with auto-upgrades the reference used as the log never needs to be explicitly upgraded and can always be written to—an interface expected of a log.

Recall that **toLabeled** is used to avoid label creep by allowing code to only temporarily raise the current label. Unfortunately, with auto-upgrades, when the current label gets raised within a **toLabeled** block, the upgrades of the flow-sensitive references remain even after the current label is restored. Thus, reading from any flow-sensitive reference after the **toLabeled** block will raise the current label to (at least) the current label at the end of the **toLabeled** block (since all references are upgraded every time the current label gets raised). This can be used to carry out a *poison pill*-like attack [18], wherein the (usually untrusted) computation executing within the **toLabeled** block will render the outer computation useless via label creep. (We note that this attack is possible in  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$  without the auto-upgrade, but requires the attacker to manually insert all the upgrades.)

To address this issue, we extend  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$  (and  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$ ) with **withRefs**<sub>FS</sub>  $v \ t$ , which takes a bag (strict heterogeneous list)  $v$  of references and a computation  $t$ , and executes  $t$  in a configuration where the flow-sensitive reference store only contains the subset of references  $v$  (and any nested references). This extension and type rule (TYPE-WITHREF), which ensures that a term cannot access a reference outside its store, are shown in Figure 11.

A bag is either empty  $\epsilon_{\overline{\tau}, \dots}$ , or it may contain a set of references of (potentially) distinct types  $\overline{v}, \dots$ . Rules (WITHREFS-CTX) and (WITHREFS-DONE) precisely define the semantics of this new primitive, where the meta-level function  $\text{addrs}(\cdot)$  converts a bag of references to a set of their corresponding addresses,  $\text{addrs}^{-1}(\cdot)$  performs the inverse conversion, and  $\ltimes$  is used to merge the stores, giving preferences to the left-hand-side store, i.e., when there is a discrepancy on a stored value between both stores, it chooses the one appearing on the left-hand-side. The function  $\text{addrs}_{\mu}^{+}(\cdot)$  computes the closure of  $\text{addrs}(\cdot)$  under store  $\mu$ , so as to include the addresses of arbitrarily-nested references. Note that if we did not include these addresses in the restricted store  $\mu'_{\text{FS}}$ , evaluation might get

$$\begin{aligned}
v &::= \dots \mid \overline{v, \dots} \mid \epsilon_{\overline{\tau, \dots}} \\
t &::= \dots \mid \overline{t, \dots} \mid \mathbf{withRefs}_{\text{FS}} \ t \ t \\
\tau &::= \dots \mid \overline{\tau, \dots} \\
E &::= \dots \mid \overline{E, t, \dots} \mid \overline{v, E, t, \dots} \mid \mathbf{withRefs}_{\text{FS}} \ E \ t \\
\\
\text{addr}_\mu(\epsilon_{\overline{\tau, \dots}}) &\triangleq \emptyset \\
\text{addr}_\mu(\overline{Ref_{\text{FS}}^{\text{TCB}} a_1, Ref_{\text{FS}}^{\text{TCB}} a_2, \dots}) &\triangleq \{a_1, a_2, \dots\} \\
\text{addr}_\mu^+(\epsilon_{\overline{\tau, \dots}}) &\triangleq \emptyset \\
\text{addr}_\mu^+(\overline{v_1, v_2, \dots}) &\triangleq \bigcup \{ \text{addr}_\mu^+(v_1), \text{addr}_\mu^+(v_2), \dots \} \\
\text{addr}_\mu^+(\overline{Ref_{\text{FS}}^{\text{TCB}} a}) &\triangleq \{a\} \cup \text{addr}_\mu^+(\mu(a)) \\
\text{addr}_\mu^+(v) &\triangleq \emptyset \\
\\
\text{WITHREFS-CTX} \\
\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \mu'_{\text{FS}} = \{a \mapsto \mu_{\text{FS}}(a) \mid a \in \text{dom } \mu_{\text{FS}} \cap (\text{addr}_{\mu_{\text{FS}}}^+(v))\} \\
\langle l_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}} \mid \mathbf{E}[t] \rangle \longrightarrow \langle l'_{\text{cur}}, \mu'_{\text{FI}}, \mu''_{\text{FS}} \mid \mathbf{E}[t'] \rangle \\
\Sigma'' = (l'_{\text{cur}}, \mu'_{\text{FI}}, \mu''_{\text{FS}} \ltimes \mu_{\text{FS}}) \quad v' = \text{addr}_{\mu''_{\text{FS}}}^{-1}(\text{dom } \mu''_{\text{FS}}) \\
\hline
\langle \Sigma \mid \mathbf{E}[\mathbf{withRefs}_{\text{FS}} \ v \ t] \rangle \longrightarrow \langle \Sigma'' \mid \mathbf{E}[\mathbf{withRefs}_{\text{FS}} \ v' \ t'] \rangle \\
\\
\text{WITHREFS-DONE} \\
\hline
\langle \Sigma \mid \mathbf{E}[\mathbf{withRefs}_{\text{FS}} \ v \ v'] \rangle \longrightarrow \langle \Sigma \mid \mathbf{E}[v'] \rangle \\
\\
\text{TYPE-WITHREF} \\
\Delta' = \{a \mapsto \Delta(a) \mid a \in \text{dom } \Delta \cap (\text{addr}_\mu(v))\} \\
\Delta', \Gamma \vdash v : \overline{Ref_{\text{FS}} \ \tau_1, \dots} \quad \Delta', \Gamma \vdash t : LIO \ \tau \\
\hline
\Delta, \Gamma \vdash \mathbf{withRefs}_{\text{FS}} \ v \ t : LIO \ \tau
\end{aligned}$$

Fig. 11: Extending  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$  and  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$  with  $\mathbf{withRefs}_{\text{FS}}$ .

stuck if the program attempted a  $\mathbf{readRef}_{\text{FS}}$  operation on a nested reference. We note that (WITHREFS-CTX) is triggered until the term under evaluation is reduced to a value, at which point (WITHREFS-DONE) is triggered, returning said value; we specify this big-step rule in terms of small-steps to facilitate the formalization of our concurrent calculus (see Section 4). Aside from the modeling of bags, the  $\mathbf{withRefs}_{\text{FS}}$  primitive is straightforward and mostly standard; indeed, the programming paradigm is similar to that already present in some mainstream languages (e.g., C++’s lambda closures require the programmer to specify the captured references). Lastly, we note that the poison pill attack can now be addressed by simply wrapping  $\mathbf{toLabeled}$  with  $\mathbf{withRefs}_{\text{FS}}$ , which prevents (untrusted) code within the  $\mathbf{toLabeled}$  block from upgrading arbitrary references.

## 4 Concurrency

In this section, we consider flow-sensitive references in the presence of concurrency (e.g., a web application in which request-handling threads share a common



log). Concretely, we extend our sequential  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  and  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$  calculi with threads and a new terminal, **forkLIO**, for dynamically creating new threads, as in the concurrent version of LIO [36]. Intuitively, this concurrent calculus  $\lambda_{\ell}^{\parallel\text{-LIO}}$  simply defines a scheduler over sequential threads, such that taking a step in the concurrent calculus amounts to taking a step in a sequential thread and context switching to a different one. For brevity, we restrict our discussion in this section to the case where the underlying sequential calculus is  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$ , since this calculus extends  $\lambda_{\ell, \text{FS}}^{\text{uo}}$ .

$$\begin{array}{c}
t ::= \dots \mid \mathbf{forkLIO} \ t \mid \text{toLabeled } \tau \ \bar{\tau} \\
\\
\text{FORKLIO} \\
\hline
\langle \Sigma \mid \mathbf{E} [\mathbf{forkLIO} \ t] \rangle \xrightarrow{\text{fork}(t)} \langle \Sigma' \mid \mathbf{E} [\mathbf{return} \ ()] \rangle \\
\\
\text{WITHREFS-OPT} \\
\frac{v = \text{addr}^{-1}((\text{addr}(v_1)) \cap (\text{addr}(v_2))) \quad \langle \Sigma \mid \mathbf{E} [\mathbf{withRefs}_{\text{FS}} \ v \ t] \rangle \longrightarrow \langle \Sigma' \mid \mathbf{E} [t'] \rangle}{\langle \Sigma \mid \mathbf{E} [\mathbf{withRefs}_{\text{FS}} \ v_1 \ (\mathbf{withRefs}_{\text{FS}} \ v_2 \ t)] \rangle \longrightarrow \langle \Sigma' \mid \mathbf{E} [t'] \rangle} \\
\\
\text{T-STEP} \\
\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \langle \Sigma \mid \mathbf{withRefs}_{\text{FS}} \ v \ t \rangle \longrightarrow \langle \Sigma' \mid t' \rangle \quad \Sigma' = (l'_{\text{cur}}, \mu'_{\text{FI}}, \mu'_{\text{FS}}) \quad v' = \text{addr}^{-1}(\text{dom } \mu'_{\text{FS}})}{\{\mu_{\text{FI}}, \mu_{\text{FS}} \mid \langle l_{\text{cur}}, v, t \rangle, k_2, \dots \} \longrightarrow \{\mu'_{\text{FI}}, \mu'_{\text{FS}} \mid k_2, \dots, \langle l'_{\text{cur}}, v', t' \rangle \}} \\
\\
\text{T-STUCK} \\
\hline
\{\mu_{\text{FI}}, \mu_{\text{FS}} \mid \langle l_{\text{cur}}, v, \uparrow \rangle, k_2, \dots \} \longrightarrow \{\mu_{\text{FI}}, \mu_{\text{FS}} \mid k_2, \dots \} \\
\\
\text{T-DONE} \\
\hline
\{\mu_{\text{FI}}, \mu_{\text{FS}} \mid \langle l_{\text{cur}}, v, v' \rangle, k_2, \dots \} \longrightarrow \{\mu_{\text{FI}}, \mu_{\text{FS}} \mid k_2, \dots \} \\
\\
\text{T-FORK} \\
\frac{\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \quad \langle \Sigma \mid \mathbf{withRefs}_{\text{FS}} \ v \ t \rangle \xrightarrow{\text{fork}(t')} \langle \Sigma' \mid t'' \rangle \quad \Sigma' = (l'_{\text{cur}}, \mu'_{\text{FI}}, \mu'_{\text{FS}}) \quad v' = \text{addr}^{-1}(\text{dom } \mu'_{\text{FS}}) \quad k_{\text{new}} = \langle l'_{\text{cur}}, v', t' \rangle}{\{\mu_{\text{FI}}, \mu_{\text{FS}} \mid \langle l_{\text{cur}}, v, t \rangle, k_2, \dots \} \longrightarrow \{\mu'_{\text{FI}}, \mu'_{\text{FS}} \mid k_2, \dots, \langle l'_{\text{cur}}, v', t'' \rangle, k_{\text{new}} \}}
\end{array}$$

Fig. 12: Semantics for  $\lambda_{\ell}^{\parallel\text{-LIO}}$ , parametric in the flow-sensitivity policy, i.e., with and without auto-upgrade.

Figure 12 shows our extended concurrent calculus,  $\lambda_{\ell}^{\parallel\text{-LIO}}$ . A concurrent program configuration has the form  $\{\mu_{\text{FI}}, \mu_{\text{FS}} \mid k_1, k_2, \dots\}$ , where  $\mu_{\text{FI}}$  and  $\mu_{\text{FS}}$  are respectively the flow-insensitive and flow-sensitive stores shared by all the threads  $k_1, k_2, \dots$  in the program. Since the memory stores are global, a thread  $k$  is simply a tuple encapsulating the current label of the thread  $l_{\text{cur}}$ , the term under evaluation  $t$ , and a bag of references  $v$  the thread may access, i.e.,  $k = \langle l_{\text{cur}}, v, t \rangle$ .

The reduction rules for concurrent programs are mostly standard. Rule (T-STEP) specifies that if the first thread in the thread pool takes a step in  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$ , the whole concurrent program takes a step, moving the thread to the end of the pool. We note that the thread term  $t$  executed with its stored current label  $l_{\text{cur}}$ , and a subset of the flow-sensitive memory store, by wrapping it in **withRefs**<sub>FS</sub>. While the use of **withRefs**<sub>FS</sub> makes the extension straightforward, one peculiarity arises: since (T-STEP) always wraps the thread term  $t$  with **withRefs**<sub>FS</sub>, if  $t$  does not reduce in one step to a value, and instead reduces to a term  $t'$ , the next time the thread is scheduled, we will superfluously wrap **withRefs**<sub>FS</sub>  $t'$  with yet another **withRefs**<sub>FS</sub>—thus preventing the thread from making progress! To address this problem, we extend the calculus with rule (WITHREFS-OPT) that collapses nested **withRefs**<sub>FS</sub> blocks.<sup>7</sup>

Rules (T-DONE) and (T-STUCK) specify that once a thread term has reduced to a value or got stuck, which is represented by  $\uparrow$ , the scheduler removes it from the thread pool and schedules the next thread.

As shown in Figure 12, to allow for dynamic thread creation, we extend  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$ 's terms with **forkLIO**, and add a new reduction rule that sends an event to the scheduler, specifying the term to execute in a new thread.<sup>8</sup> Rule (T-FORK) describes the corresponding scheduler rule, triggered when a *fork* ( $t'$ ) event is received. Here, we create a new thread  $k_{\text{new}}$  whose current label  $l'_{\text{cur}}$  and partition of the store, i.e., bag of references  $v'$ , is the same as that of the parent thread; the term evaluated in the newly created thread is provided in the event:  $t'$ . Subsequently, we add the new thread to the thread pool.

The final modification in extending  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$  to  $\lambda_{\ell}^{\parallel\text{-LIO}}$  is the removal of **toLabeled** from the underlying calculus. As described in [36], we must remove **toLabeled** to guarantee termination-sensitive non-interference. Importantly, however, **forkLIO** with synchronization primitives (e.g., flow-insensitive labeled MVars, as discussed in [36]) can be used to provide functionality equivalent to that of **toLabeled**. Due to space constraints we omit synchronization primitives from  $\lambda_{\ell}^{\parallel\text{-LIO}}$ ; we only remark that extending  $\lambda_{\ell}^{\parallel\text{-LIO}}$  to provide flow-sensitive labeled MVars follows in a straightforward way.

Since the flow-sensitive attack in Figure 7 relied on **toLabeled** to restore the current label, a natural question, given that we remove **toLabeled**, is whether we can use the naive flow-sensitive reference semantics of Section 3 for concurrent LIO. As shown by the attack code in Figure 13, in which we use **forkLIO** instead of **toLabeled** to address a potential label creep, the fundamental problem remains: the label on the reference label is not protected! This precisely motivated our principled approach of extending  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$  to a concurrent setting as opposed to extending concurrent LIO with flow-sensitive references.

<sup>7</sup> This change also requires modifying (WITHREFS-CTX) to not be triggered when the term being evaluated is a **withRefs**<sub>FS</sub> term.

<sup>8</sup> In fact, the reduction rule for  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{LIO}}$  must be changed to account for events. However, since *fork* is the only event in our system, we treat  $\rightarrow$  as implicitly carrying an empty event.

```

leakRef :: RefTCB Bool → LIO Bool
leakRef href = do
  tmp ← newRef L ()
  forkLIO $ do h ← readRef href
               when h $ writeRef tmp ()
  delay
  return $ labelOf tmp ≡ H

```

Fig. 13: Attack on concurrent LIO with naive flow-sensitive reference extension.

## 5 Formal results

In this section, we show that our flow-sensitive enforcement can be embedded into the flow-insensitive version of LIO. Additionally, we provide security guarantees in terms of non-interference definitions by reusing previous results on LIO.

### 5.1 Embedding into $\lambda_\ell^{\text{uo}}$

Every flow-sensitive reference with label  $l_d$  created in a context where the current label is  $l_o$  (and thus stored in  $\mu_{\text{FS}}$  as  $Lb^{\text{TCB}} l_o (Lb^{\text{TCB}} l_d t)$ ), can be represented by a flow-insensitive reference with label  $l_o$ , whose contents are another flow-insensitive reference containing  $t$  and labeled  $l_d$ .

Figure 14 gives our encoding of the flow-sensitive reference operations in terms of flow-insensitive references. For a given store  $\Sigma$ , we define the  $\llbracket - \rrbracket_{\text{FI}}^\Sigma$  function, which given a term  $t$  in  $\lambda_{\ell, \text{FS}}^{\text{uo}}$ , produces a term  $\llbracket t \rrbracket_{\text{FI}}^\Sigma$  in  $\lambda_\ell^{\text{uo}}$ , expanding the definitions of flow-sensitive operations in terms of flow-insensitive ones. This function is applied homomorphically in all other cases. We use the *WrapRef* constructor to mark the flow-insensitive references that are being used to represent flow-sensitive ones, so as to distinguish them from other flow-insensitive references. The functions *wrap* and *unwrap* are used to add and remove this boundary encoding. In the embedding of **writeRef**<sub>FS</sub>, we use **toLabeled** to make any changes to the current label (possibly caused by reading the outer reference) local to this operation. Inside **toLabeled**, the code fetches the inner reference (**readRef**<sub>FI</sub>), and then performs the actual write of the new value. If this fails, the computation diverges, but, importantly, the current label was raised (with **readRef**<sub>FI</sub>) to reflect the fact that the label on the label of the reference was observed. The embedding of **upgrade**<sub>FS</sub> relies on flow-insensitive primitives to implement the **upgrade** operation. As in **writeRef**<sub>FS</sub>, a **toLabeled** block is used to delimit the taint on the current label. Inside the block, the code fetches the inner reference (**readRef**<sub>FI</sub>), which taints the current label with  $l'$ , and makes a new reference  $n$  (**newRef**<sub>FI</sub>) with the upgraded label  $(l_{\text{cur}} \sqcup (l \sqcup \text{labelOf } i))$  and an undefined value ( $\perp$ ). Observe that the operation for creating the reference always succeeds since its label is above the current label, i.e.,  $l_{\text{cur}}$ . Then,

**copyRef** is used to copy the value of the original inner reference into the new one,  $n$ . As before, this action always succeeds because the label of the reference bound to  $i$  (**labelOf**  $i$ ) is below the label of the new reference  $n$ . Finally, the reference  $n$  is stored in place of the original inner reference using **writeRef**. Importantly, this instruction only succeeds when the current label at the time of writing, i.e.,  $lc \sqcup l'$  in Figure 14, is below or equal to  $l'$  (the label of the outer reference), i.e.,  $lc \sqcup l' \sqsubseteq l'$ . This restriction holds when the current label at the time of upgrade, i.e.,  $lc$ , is below or equal to  $l'$ —effectively encoding the non-sensitive upgrade policy for label changes. The embedding of **downgrade**<sub>FS</sub> follows similarly, except that the label of the new reference is computed using  $\sqcap$  instead of  $\sqcup$  (to achieve the downgrade), and the **copyRef** step is omitted, since the original value must be destroyed. We remark that the mapping mimics the behavior described by the rules in Figure 8.

We extend this definition naturally to convert  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$  environments into  $\lambda_{\ell}^{\text{LIO}}$  environments, by having  $\llbracket (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}}) \rrbracket_{\text{FI}} \triangleq (l_{\text{cur}}, \mu'_{\text{FI}})$  where  $\mu'_{\text{FI}}$  is obtained by extending  $\mu_{\text{FI}}$  with the pair of bindings  $a_i \mapsto Lb^{\text{TCB}} l_o (Ref^{\text{TCB}}_{\text{FI}} l_d b_i)$ ,  $b_i \mapsto (Lb^{\text{TCB}} l_d v)$  (with  $b_i$  being a fresh name) for each binding of the form  $a_i \mapsto Lb^{\text{TCB}} l_o (Lb^{\text{TCB}} l_d v_i)$  in  $\mu_{\text{FS}}$ . Note that the domains of  $\mu_{\text{FI}}$  and  $\mu_{\text{FS}}$  are disjoint because the fresh( $\cdot$ ) predicate that we use in the semantics is assumed to produce globally unique addresses.

In order to prove that our implementation is correct with respect to the semantics, we show that, if we take a program with flow-sensitive operations, and expand those operations, replacing them by the code in Figure 14, then its behavior corresponds with the flow-sensitive semantics.

**Theorem 1 (Embedding  $\lambda_{\ell, \text{fs}}^{\text{LIO}}$  in  $\lambda_{\ell}^{\text{LIO}}$ ).** *Let  $t$  be a well-typed term in  $\lambda_{\ell, \text{fs}}^{\text{LIO}}$ . Then if  $\langle \Sigma | t \rangle \longrightarrow^* \langle \Sigma' | v \rangle$ , we have  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket t \rrbracket_{\text{FI}}^{\Sigma} \rangle \longrightarrow^* \langle \llbracket \Sigma' \rrbracket_{\text{FI}} | \llbracket v \rrbracket_{\text{FI}}^{\Sigma} \rangle$ , and if  $\langle \Sigma | t \rangle \longrightarrow^* \langle \Sigma' | \uparrow \rangle$ , then  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket t \rrbracket_{\text{FI}}^{\Sigma} \rangle \longrightarrow^* \langle \llbracket \Sigma' \rrbracket_{\text{FI}} | \uparrow \rangle$ .*

*Proof.* See Appendix C.

While straight forward, this theorem highlights an important result: in floating label systems, flow-sensitive references can be encoded in a calculus with flow-insensitive references and explicitly labeled values.

## 5.2 Security guarantees for $\lambda_{\ell, \text{fs}}^{\text{LIO}}$ , $\lambda_{\ell, \text{fs}+\text{au}}^{\text{LIO}}$ and $\lambda_{\ell}^{\text{LIO}}$

From previous results [35], we know that LIO satisfies termination-insensitive non-interference (TINI) in the sequential setting, and termination-sensitive non-interference (TSNI) in the concurrent setting. By using the embedding theorem we can extend these results for LIO with flow-sensitive references.

For completeness, we now present our non-interference theorems, as straightforward applications of the theorems in previous work. Our security results rely on the notion of  $l$ -equivalence for terms and configurations, which captures the idea of terms that cannot be distinguished by an attacker which can observe data at level  $l$ . A pair of terms  $t_1, t_2$  is said to be  $l$ -equivalent (written  $t_1 \approx_l t_2$ ) if,

$$\begin{aligned}
\text{wrap } r &\triangleq \text{WrapRef } r \\
\text{unwrap } (\text{WrapRef } r) &\triangleq r \\
\llbracket \text{Ref}_{\text{FS}}^{\text{TCB}} r \rrbracket_{\text{FI}}^{(l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}})} &\triangleq \text{wrap } (\text{Ref}_{\text{FI}}^{\text{TCB}} (\text{labelOf}_{\text{FI}} \mu_{\text{FS}} (r)) r) \\
\llbracket \text{newRef}_{\text{FS}} \rrbracket_{\text{FI}}^{\Sigma} &\triangleq \lambda l \ t. \text{do} \\
&\quad i \leftarrow \text{newRef}_{\text{FI}} l \ t \\
&\quad l_{\text{cur}} \leftarrow \text{getLabel} \\
&\quad o \leftarrow \text{newRef}_{\text{FI}} l_{\text{cur}} i \\
&\quad \text{return } (\text{wrap } o) \\
\llbracket \text{readRef}_{\text{FS}} \rrbracket_{\text{FI}}^{\Sigma} &\triangleq \lambda r. \text{readRef}_{\text{FI}} (\text{unwrap } r) \gg \text{readRef}_{\text{FI}} \\
\llbracket \text{writeRef}_{\text{FS}} \rrbracket_{\text{FI}}^{\Sigma} &\triangleq \lambda r \ t. \text{let } o = \text{unwrap } r \text{ in do} \\
&\quad l_{\text{cur}} \leftarrow \text{getLabel} \\
&\quad \text{toLabeled } (l_{\text{cur}} \sqcup (\text{labelOf } o)) \$ \text{do} \\
&\quad \quad i \leftarrow \text{readRef}_{\text{FI}} o \\
&\quad \quad \text{writeRef}_{\text{FI}} i \ t \\
\llbracket \text{labelOf}_{\text{FS}} \rrbracket_{\text{FI}}^{\Sigma} &\triangleq \lambda r. \\
&\quad \text{readRef}_{\text{FI}} (\text{unwrap } r) \gg \text{return.labelOf}_{\text{FI}} \\
\llbracket \text{upgrade}_{\text{FS}} \rrbracket_{\text{FI}}^{\Sigma} &\triangleq \lambda r \ l. \text{let } o = \text{unwrap } r \\
&\quad \quad l' = \text{labelOf } o \text{ in do} \\
&\quad \quad lc \leftarrow \text{getLabel} \\
&\quad \quad \text{toLabeled } (lc \sqcup l') \$ \text{do} \\
&\quad \quad \quad i \leftarrow \text{readRef}_{\text{FI}} o \\
&\quad \quad \quad l_{\text{cur}} \leftarrow \text{getLabel} \\
&\quad \quad \quad n \leftarrow \text{newRef}_{\text{FI}} (l_{\text{cur}} \sqcup (l \sqcup \text{labelOf } i)) \perp \\
&\quad \quad \quad \text{copyRef } i \ n \\
&\quad \quad \quad \text{writeRef}_{\text{FI}} o \ n \\
\llbracket \text{downgrade}_{\text{FS}} \rrbracket_{\text{FI}}^{\Sigma} &\triangleq \lambda r \ l. \text{let } o = \text{unwrap } r \\
&\quad \quad l' = \text{labelOf } o \text{ in do} \\
&\quad \quad lc \leftarrow \text{getLabel} \\
&\quad \quad \text{toLabeled } (lc \sqcup l') \$ \text{do} \\
&\quad \quad \quad i \leftarrow \text{readRef}_{\text{FI}} o \\
&\quad \quad \quad l_{\text{cur}} \leftarrow \text{getLabel} \\
&\quad \quad \quad n \leftarrow \text{newRef}_{\text{FI}} (l_{\text{cur}} \sqcup (l \sqcap \text{labelOf } i)) \perp \\
&\quad \quad \quad \text{writeRef}_{\text{FI}} o \ n \\
\llbracket \text{withRefs}_{\text{FS}} v \ t \rrbracket_{\text{FI}}^{(l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}})} &\triangleq \llbracket t \rrbracket_{\text{FI}}^{(l_{\text{cur}}, \mu_{\text{FI}}, \mu'_{\text{FS}})} \\
\text{where} & \\
\mu'_{\text{FS}} &= \{ a \mapsto \mu_{\text{FS}}(a) \mid a \in \text{dom } \mu_{\text{FS}} \cap (\text{addrs}_{\mu_{\text{FS}}}^+(v)) \}
\end{aligned}$$

Fig. 14: Implementation mapping for flow-sensitive references. For all other terms, the function is applied homomorphically.

after erasing all the information more sensitive than  $l$  from  $t_1$  and  $t_2$ , we obtain syntactically equivalent terms. This definition extends naturally to configurations.

Intuitively, non-interference means that an attacker at level  $l$  cannot distinguish among different runs of a program with  $l$ -equivalent initial configurations.

**Theorem 2 (TINI for  $\lambda_{\ell, \text{fs}}^{\text{uo}}$ ).** *Consider two well-typed terms  $t_1$  and  $t_2$  in  $\lambda_{\ell, \text{fs}}^{\text{uo}}$  which do not contain any  $\cdot^{\text{TCB}}$  syntax nodes, such that  $t_1 \approx_l t_2$ , where  $l$  is the attacker observation level. Let  $\Sigma$  be an initial environment, and let*

$$\langle \Sigma | t_1 \rangle \longrightarrow^* \langle \Sigma_1 | v_1 \rangle \text{ and } \langle \Sigma | t_2 \rangle \longrightarrow^* \langle \Sigma_2 | v_2 \rangle$$

*Then, we have that  $\langle \Sigma_1 | v_1 \rangle \approx_l \langle \Sigma_2 | v_2 \rangle$ .*

*Proof.* By expanding all the flow-sensitive operations in  $t_1$  and  $t_2$  using their definition given in Figure 14, we get terms in  $\lambda_{\ell}^{\text{uo}}$ , which by Theorem 1 has equivalent semantics. Therefore, the result follows from the  $\lambda_{\ell}^{\text{uo}}$  TINI result of [35].

**Corollary 1 (TINI for  $\lambda_{\ell, \text{fs}+\text{au}}^{\text{uo}}$ ).** *The previous non-interference result can be easily extended to  $\lambda_{\ell, \text{fs}+\text{au}}^{\text{uo}}$ . In  $\lambda_{\ell, \text{fs}+\text{au}}^{\text{uo}}$ , the **unlabel** operation triggers the automatic upgrades mechanism, which performs the **upgrade** operation for every flow-sensitive reference in scope before actually raising the current label. Regardless of how **unlabel** is used, we note that the resulting term (after inserting the necessary **upgrades**), is just an  $\lambda_{\ell, \text{fs}}^{\text{uo}}$  term. Therefore, the main TINI result for  $\lambda_{\ell, \text{fs}}^{\text{uo}}$  applies.*

For the concurrent result, we need a supporting lemma which states that the current label is always at least as sensitive as the label of every reference in scope. Formally,

**Lemma 1.** *Let  $t$  be a well-typed term in  $\lambda_{\ell}^{\parallel\text{-LIO}}$ ,  $\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}})$  be an initial environment, and  $a$  be the address of a flow-sensitive reference  $r$  in  $\Sigma$ , where  $\mu_{\text{FS}}(a) = \text{Lb}^{\text{TCB}} l_o (\text{Lb}^{\text{TCB}} l_d v)$ . Then, if  $\langle \Sigma | t \rangle \longrightarrow^* \langle l'_{\text{cur}}, \mu'_{\text{FI}}, \mu'_{\text{FS}} | t' \rangle$ , we have that  $l_o \sqsubseteq l'_{\text{cur}}$ .*

*Proof.* Note that the result holds immediately after creating  $r$ , since the current label is the label on the label of  $r$ , i.e.,  $l_o = l_{\text{cur}}$ . It is easy to show that  $l_o$  is immutable, since there are no reduction rules that modify it. Moreover, given that the current label is monotonic, the only way in which  $l_o \sqsubseteq l_{\text{cur}}$  can cease to hold is if  $r$  is accessed from a different thread. But in order to pass  $r$  to a different thread, a labeled object must be used as intermediary, and the label of such object would have to be at least  $l_{\text{cur}}$ , the current label in the thread that created  $r$ . As a result, if we were to pass  $r$  to another thread in this way, then the target thread would also have to be tainted by  $l_{\text{cur}}$ , and the result would still hold.

We now prove our non-interference theorem for  $\lambda_{\ell}^{\parallel\text{-LIO}}$ . This result is stronger than TINI, since it implies that there can be no termination or internal timing leaks.

**Theorem 3 (TSNI for  $\lambda_\ell^{\text{LIO}}$ ).** Consider two well-typed terms  $t_1$  and  $t_2$  in  $\lambda_\ell^{\text{LIO}}$  which do not contain any  $\cdot^{\text{TCB}}$  syntax nodes, such that  $t_1 \approx_l t_2$ , where  $l$  is the attacker observation level. Let  $\Sigma = (l_{\text{cur}}, \mu_{\text{FI}}, \mu_{\text{FS}})$  be an initial environment, and let

$$\{\mu_{\text{FI}}, \mu_{\text{FS}} \mid \langle l_{\text{cur}}, \text{addrs}^{-1}(\text{dom } \mu_{\text{FS}}), t_1 \rangle\} \longrightarrow^* M_1$$

Then, there exists some configuration  $M_2$  such that

$$\{\mu_{\text{FI}}, \mu_{\text{FS}} \mid \langle l_{\text{cur}}, \text{addrs}^{-1}(\text{dom } \mu_{\text{FS}}), t_2 \rangle\} \longrightarrow^* M_2$$

and  $M_1 \approx_l M_2$ .

*Proof.* From Lemma 1 and looking at the embeddings of **writeRef**<sub>FS</sub> and **upgrade**<sub>FS</sub>, we note that the first **readRef**<sub>FI</sub> operation in each **toLabeled** block will be trying to raise the current label to  $l$ . However, since  $l \sqsubseteq l_{\text{cur}}$ , these operations will never effectively raise the current label. This means that using **toLabeled** is not necessary to preserve the semantics, because there is no need to restore the current label afterwards. As a result, and after removing **toLabeled** in these two cases, we note that the embedding produces valid concurrent  $\lambda_\ell^{\text{LIO}}$  terms (which does not have **toLabeled**).

Finally, by expanding all the flow-sensitive operations in  $t_1$  and  $t_2$  using their definition given in Figure 14, we get terms in concurrent  $\lambda_\ell^{\text{LIO}}$ . Therefore, the result follows from the termination-sensitive non-interference of concurrent  $\lambda_\ell^{\text{LIO}}$  [36].

The detailed proofs for the results in this section can be found in Appendix C. We lastly remark a limitation: while we preserve non-interference when embedding the flow-sensitive calculus in the original LIO, our embedding includes no synchronization to ensure atomicity of the flow-sensitive operations, so certain interleaving that break semantic equivalence are possible.

### 5.3 Permissiveness

In Section 7 we compare the permissiveness of our system with previous flow-sensitive IFC systems. Here, we solely remark that the above results imply that our flow-sensitive calculus is as permissive as flow-insensitive LIO. In particular, any flow-insensitive LIO program can be trivially converted to a flow-sensitive LIO program (without auto-upgrades) by using flow-sensitive references instead of flow-insensitive ones. Since these references would never be upgraded, they will behave just like their flow-insensitive counterparts. This means that all existing LIO programs can be run in our flow-sensitive monitor. This includes Hails [14], a web framework using LIO, on top of which a number of applications have been built (e.g., GitStar<sup>9</sup>, a code-hosting web platform, LearnByHacking<sup>10</sup>, a blog/tutorial platform similar to School of Haskell, and LambdaChair [38], an EasyChair-like conference review system).

<sup>9</sup> [www.gitstar.com](http://www.gitstar.com)

<sup>10</sup> [www.learnbyhacking.org](http://www.learnbyhacking.org)

## 6 Comparison with other policies for label change

In this section, we compare our enforcement mechanism with two policies for label change: *no-sensitive-upgrade* (originally proposed by Zdancewic [43]) and *permissive-upgrade*, a more permissive version of the former by Austin and Flanagan [2, 3]. We introduce a simple imperative language to simplify our comparison with the languages implementing the aforementioned policies. This simple language has variables, **if**-statements, a **skip** command that does nothing, and an **output** command that is used to produce public outputs. This language is easily implemented in  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  as syntactic sugar. For example, a conditional statement **if**  $C$ ;  $A$  **else**  $B$  is desugared to **toLabeled**  $\mathbf{H}$  (**do**  $b \leftarrow \text{unlabel } C$ ; **if**  $b$  **then**  $A$  **else**  $B$ ). (Here, **toLabeled** is used to ensure that the current label is restored after leaving the **if**-statement.)

### 6.1 No-sensitive-upgrade

The no-sensitive-upgrade discipline stops execution on any attempt to change the label of a public variable inside a secret context. Our  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  calculus essentially implements this discipline as well—see (UPGRADE<sub>REF</sub>) in Figure 8. The original presentation of no-sensitive-upgrade allows for variables with a secret label to be downgraded, as long as the original contents are discarded. Our  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  calculus similarly allows for this with the **downgrade** operation. Our approach differs in also allowing code to explicitly upgrade a variable before entering a secret context, permissively allowing writes to originally-public variables in secret contexts.

### 6.2 Permissive-upgrade

The permissive-upgrade policy differs from no-sensitive-upgrade in allowing code to change the label of a public variables in secret contexts, but subsequently disallowing branches on such permissively-upgraded, or “marked,” variables. When upgrading a public variable in a secret context, the security label of the variable is changed to  $\mathbf{P}$  where  $\mathbf{L} \sqsubseteq \mathbf{H} \sqsubseteq \mathbf{P}$ .

```

upgrade  $x$   $\mathbf{H}$ 
if  $h$ 
   $x := \text{True}$ 
if  $x$ 
  skip

```

Fig. 15: A secure program that  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  accepts.

In general, the permissiveness of our approach is incomparable with that of permissive-upgrade. For example, without the **upgrade** operation,  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  is as expressive as no-sensitive-upgrade, and thus less permissive than permissive-upgrade. But, with **upgrade** we can write programs in  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  that would be rejected by a permissive-upgrade monitor. Figure 15 shows an example of one such case. In the example, the **upgrade** operation is used to ensure that reference  $x$ , which would be marked  $\mathbf{P}$  by permissive-upgrade, ends up as  $\mathbf{H}$  in all runs; without the **upgrade**, a permissive-upgrade monitor would reject the branch on  $x$ . By inserting



**upgrade** operations in the “right” places, our approach can become more flexible than permissive-upgrade.

```

if  $h$ 
   $x := \text{True}$ 
if  $x$ 
  skip

```

Fig. 16: A secure program that permissive-upgrade rejects and  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$  accepts.

Of course, the challenge lies in upgrading references in a permissive fashion. And automatically upgrading references whenever the current label is raised is not necessarily more permissive than a permissive-upgrade monitor. Indeed, the permissiveness of permissive-upgrade and  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$  are also incomparable.

Figure 16 shows a program that is rejected by permissive-upgrade but accepted by  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$ . With permissive-upgrade, the first branch on  $h$  upgrades  $x$  to **P**, which leads to a failure when subsequently trying to branch  $x$ . With  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$ , on the other hand, reference  $x$  would be upgraded to **H**, permissively allowing the second branch.

Conversely, Figure 17 shows a secure program that is accepted by permissive-upgrade but rejected by  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$ . In this program, there are two variables in scope:  $x$  and  $y$ , both initially public. In the first secret conditional block we assign to  $y$ , which with permissive-upgrade only upgrades variable  $y$  to **P**;  $x$  remains **L** and thus the second branch is executed, producing a public output. With  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$ , however, unlabeled  $h$  (an operation implicit in the first conditional, which inspects  $h$ ) auto-upgrades both  $x$  and  $y$  to **H**. As a result, the current label at the point of the **output** is **H**, causing a failure.

One way to address the permissiveness issues of  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$  is by using **withRefs<sub>FS</sub>** to delimit the scope of the upgrades. Figure 18 shows a modified version of the previous example, where  $y$  is explicitly marked as the only variable that should be upgraded in the first branch. As a consequence,  $x$  does not get upgraded and the program does not fail—the **output** operation is allowed. More generally, if there is enough static information to guide the use of **withRefs<sub>FS</sub>**, we believe that  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$  could match (or exceed) the permissiveness of permissive-upgrade.

```

 $x, y := \text{True}$ 
if  $h$ 
   $y := \text{False}$ 
if  $x$ 
  output (1)

```

Fig. 17: A secure program that permissive-upgrade accepts and  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$  rejects.

```

 $x, y := \text{False}$ 
withRefsFS ( $y$ ) {
  if  $h$ 
     $y := \text{True}$ 
}
if  $x$ 
  output (1)

```

Fig. 18: A secure program that permissive-upgrade and  $\lambda_{\ell, \text{FS}+\text{AU}}^{\text{uo}}$  with **withRefs<sub>FS</sub>** accept.

## 7 Related work

The *label on the label* could be seen as a fixed label that dictates which principals can read or modify the policy (inner label) of a flow-sensitive entity. In a different setting, trust management frameworks have explored this idea [4], where role-based rules are labeled to restrict the view on policies—the mere presence of certain policies could become inappropriate conduits of information.

Several authors propose an *existence security label* to remove leaks due to the termination covert channel [29, 30] or certain behaviors in dynamic nested data structures [16, 32]. While the existence security label and the label on the label are structurally isomorphic, they are used for different purposes and in different scenarios, e.g., inspecting labels is not allowed in [16, 29, 30, 32].

Hunt and Sands [19] show the equivalence (modulo code transformation) between flow-sensitive and flow-insensitive type-systems. Russo and Sabelfeld [31] formally pin down the challenge of mutable labels when using purely dynamic monitors. They prove that monitors require static analysis in order to be more permissive than traditional flow-sensitive type-systems. Austin and Flanagan propose a *privatization* operation to boost the permissiveness of permissive-upgrade. This technique has been recently generalized to arbitrary lattices [5]. Moreover, the privatization operation can only enforce non-interference when outputs are suppressed after branching on a marked flow-sensitive reference. Unfortunately, none of the mentioned work consider flow-sensitive in the presence of concurrency. In fact, the notion of permissive-upgrade does not easily generalize to the concurrent setting, as this would require tracking occurrences of branches across threads.

Recently, Hritcu et al. [18] propose a floating-label system called Breeze. Like LIO, Breeze allows changes in the current context label (i.e., *pc*) and only considered values with flow-insensitive labels. Given the design similarities with LIO [35], we believe that our results could be easily adapted to Breeze.

Hedin et al. [17] recently developed JSFlow, an IFC flow-sensitive monitor for JavaScript. The monitor uses the no-sensitive-upgrade label changing policy. To overcome some of the restrictions imposed by this discipline, the primitive **upgrade** is introduced to explicitly change labels. Our upgrade operation resembles that proposed by Hedin et al. Moreover, the extension to **unlabel** as described Section 3 can be seen as an automatic application of **upgrade** every time that the current label gets raised. Using testing, Birgisson et al. [6] automatically insert **upgrade** instructions to boost the permissiveness of no-sensitive-upgrade. We further extend this concept of (automatic) **upgrades** to a concurrent setting.

The Operating System IFC community has also treated the mutable label problem in the presence of purely dynamic monitors. Specifically, IFC OSes such as Asbestos [10], HiStar [44], and Flume [21] distinguish between subjects (processes), and objects (files, sockets, etc.) such that the security labels for objects are immutable, while subject labels change according to the sensitivity of data being read. As in language-based IFC systems, changing the label of subjects and object can become a covert channel, if not handled appropriated.

Hence, HiStar and Flume require that the label of a subject be done explicitly by the subject code. Asbestos, on the other hand, allowed (unsafe) changes to labels as the result of receiving messages under specific and safe conditions. Our work extends these concepts with a level of indirection to allow for changes in object labels.

Coarse-grained IFC enforcements, similar to the ones found in IFC OS work, have been applied to web browsers. e.g, BFlow [42] and COWL [39] track the flow of information at the granularity of protection zones and context, respectively. Both can be understood as tracking IFC at the iframe-level granularity. As in LIO, an iframe’s label, i.e., a subject’s label, must be explicitly updated. While our techniques can be applied to COWL, BFlow does not have fine-grained labeled objects; hence the flow-sensitivity result is only applicable at the protection zone level. By taking a more fine grained approach, the DOM-tree could be thought of as being composed of flow-sensitive objects, whose security labels change according to the dynamic behavior of the web page, as done in [32].

Hoare-like logics for IFC are often flow-sensitive [e.g. 1, 28]. Different from dynamic approaches, these logics have the ability to observe all the execution paths and safely approximate label changes. As a result, no leaks due to label changes are present in provably secure programs. Le Guernic et al. [22, 23] combine dynamic and static checks in a flow-sensitive execution monitor. For a flow-sensitive type-system, Foster et al. [13] propose a **restrict** primitive that limits the use of variables’ aliases in a block of code. Our **withRefs<sub>FS</sub>** is similar to **restrict** in being used to increase the permissiveness of the analysis.

## 8 Conclusions

We presented an extension of LIO with flow-sensitive references. As in previous flow-sensitive work, our approach allows secure label changes using **upgrade** and **downgrade** operations, as a way to boost the permissiveness of the IFC system, i.e., **upgrade** can be used to allow for the encoding of programs that would otherwise be rejected by the IFC monitor. Since manually inserting **upgrade** operations can be cumbersome, we extend the calculus to automatically insert upgrades whenever the current label is raised, while still giving programmers fine-grained control over which references untrusted code can upgrade. Importantly, our approach extends to a concurrent setting. To the best of our knowledge, this is the first work to address the problem of flow-sensitive label changes for a concurrent, dynamic IFC language. A further insight of this work was to show that, by leveraging nested labeled objects, both the sequential and concurrent calculi with flow-sensitive references can be encoded using only flow-insensitive constructs. As a consequence, our soundness proof can be reduced to an invocation of previous results for LIO.

## Acknowledgments

We thank our colleagues in the ProSec group at Chalmers, Stefan Heule, David Mazières, and Edward Z. Yang for the useful discussions. We thank the anonymous reviewers for constructive feedback on an earlier version of this work. This work was funded by DARPA CRASH under contract #N66001-10-2-4088, the Swedish research agency VR, and a grant from Mozilla. Deian Stefan was supported by the DoD through the NDSEG Fellowship Program.

## Bibliography

- [1] T. Amtoft, S. Bandhakavi, and A. Banerjee. A logic for information flow in object-oriented programs. In *Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '06. ACM, 2006.
- [2] T. H. Austin and C. Flanagan. Efficient Purely-Dynamic Information Flow Analysis. In *Proc. ACM Workshop on Programming Languages and Analysis for Security (PLAS)*, June 2009.
- [3] T. H. Austin and C. Flanagan. Permissive dynamic information flow analysis. In *Proceedings of the 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, PLAS '10. ACM, 2010.
- [4] S. Bandhakavi, W. Winsborough, and M. Winslett. A trust management approach for flexible policy management in security-typed languages. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, June 2008.
- [5] A. Bichhawat, V. Rajani, D. Garg, and C. Hammer. Generalizing permissive-upgrade in dynamic information flow analysis. In *Proceedings of the Ninth Workshop on Programming Languages and Analysis for Security*, PLAS'14. ACM, 2014.
- [6] A. Birgisson, D. Hedin, and A. Sabelfeld. Boosting the permissiveness of dynamic information-flow tracking by testing. In *Proc. European Symp. on Research in Computer Security*, 2012.
- [7] P. Buiras, D. Stefan, and A. Russo. On Dynamic Flow-sensitive Floating-Label Systems. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE, July 2014.
- [8] W. Cheng, D. R. Ports, D. Schultz, V. Popic, A. Blankstein, J. Cowling, D. Curtis, L. Shriram, and B. Liskov. Abstractions for usable information flow control in Aeolus. In *Proceedings of the 2012 USENIX Annual Technical Conference*, 2012.
- [9] W. De Groef, D. Devriese, N. Nikiforakis, and F. Piessens. FlowFox: a web browser with flexible and precise information flow control. In *Proceedings of the 2012 ACM conference on Computer and communications security*, CCS '12. ACM, 2012.
- [10] P. Efstathiopoulos, M. Krohn, S. VanDeBogart, C. Frey, D. Ziegler, E. Kohler, D. Mazières, F. Kaashoek, and R. Morris. Labels and event

- processes in the Asbestos operating system. In *Proc. of the twentieth ACM symp. on Operating systems principles*, SOSP '05. ACM, 2005.
- [11] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. TaintDroid: An Information-flow Tracking System for Realtime Privacy Monitoring on Smartphones. In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation*, OSDI'10. USENIX Association, 2010.
  - [12] M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical computer science*, 103(2):235–271, 1992.
  - [13] J. S. Foster, T. Terauchi, and A. Aiken. Flow-sensitive type qualifiers. In *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation*, PLDI '02. ACM, 2002.
  - [14] D. B. Giffin, A. Levy, D. Stefan, D. Terei, D. Mazières, J. Mitchell, and A. Russo. Hails: Protecting data privacy in untrusted web applications. In *Proc. of the 10th Symposium on Operating Systems Design and Implementation*, October 2012.
  - [15] J. A. Goguen and J. Meseguer. Security Policies and Security Models. In *Proc. IEEE Symp. on Security and Privacy*, pages 11–20, Apr. 1982.
  - [16] D. Hedin and A. Sabelfeld. Information-Flow Security for a Core of JavaScript. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, 2012.
  - [17] D. Hedin, A. Birgisson, L. Bello, and A. Sabelfeld. JSFlow: Tracking information flow in JavaScript and its APIs. In *Proc. ACM Symposium on Applied Computing (SAC)*. ACM, Mar. 2014.
  - [18] C. Hritcu, M. Greenberg, B. Karel, B. C. Pierce, and G. Morrisett. All Your IFCEException Are Belong to Us. *2012 IEEE Symposium on Security and Privacy*, 0, 2013.
  - [19] S. Hunt and D. Sands. On flow-sensitive security types. In *Conference record of the 33rd ACM SIGPLAN-SIGACT Symp. on Principles of programming languages*, POPL '06, pages 79–90. ACM, 2006.
  - [20] L. Jia, J. Aljuraidan, E. Fragkaki, L. Bauer, M. Stroucken, K. Fukushima, S. Kiyomoto, and Y. Miyake. Run-time enforcement of information-flow properties on Android (extended abstract). In *Computer Security—ESORICS 2013: 18th European Symposium on Research in Computer Security*. Springer, Sept. 2013.
  - [21] M. Krohn, A. Yip, M. Brodsky, N. Cliffer, M. F. Kaashoek, E. Kohler, and R. Morris. Information flow control for standard OS abstractions. In *Proc. of the 21st Symp. on Operating Systems Principles*, October 2007.
  - [22] G. Le Guernic. Automaton-based Confidentiality Monitoring of Concurrent Programs. In *Proc. of the 20th IEEE Computer Security Foundations Symposium*, CSF '07. IEEE Computer Society, 2007.
  - [23] G. Le Guernic, A. Banerjee, T. Jensen, and D. A. Schmidt. Automata-based confidentiality monitoring. In *Proceedings of the 11th Asian Computing Science Conference on Advances in Computer Science: Secure Software and Related Issues*, ASIAN'06. Springer-Verlag, 2006.

- [24] E. Moggi. Notions of computation and monads. *Information and Computation*, 93(1):55–92, 1991.
- [25] A. C. Myers and B. Liskov. A decentralized model for information flow control. In *Proc. of the 16th ACM Symp. on Operating Systems Principles*, pages 129–142, 1997.
- [26] A. C. Myers and B. Liskov. Protecting privacy using the decentralized label model. *ACM Trans. on Computer Systems*, 9(4):410–442, October 2000.
- [27] A. C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif: Java Information Flow. Software release. Located at <http://www.cs.cornell.edu/jif>, July 2001.
- [28] A. Nanevski, A. Banerjee, and D. Garg. Verification of information flow and access control policies with dependent types. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, SP '11. IEEE Computer Society, 2011.
- [29] W. Rafnsson and A. Sabelfeld. Secure multi-execution: fine-grained, declassification-aware, and transparent. In *Computer Security Foundations Symposium (CSF), 2013 IEEE 26th*, pages 33–48. IEEE, 2013.
- [30] W. Rafnsson, D. Hedin, and A. Sabelfeld. Securing Interactive Programs. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, 2012.
- [31] A. Russo and A. Sabelfeld. Dynamic vs. Static Flow-Sensitive Security Analysis. In *Proc. of the 2010 23rd IEEE Computer Security Foundations Symp.*, CSF '10, pages 186–199. IEEE Computer Society, 2010.
- [32] A. Russo, A. Sabelfeld, and A. Chudnov. Tracking information flow in dynamic tree structures. In *Proceedings of the 14th European Conference on Research in Computer Security*, ESORICS'09. Springer-Verlag, 2009.
- [33] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1), January 2003.
- [34] V. Simonet. The Flow Caml System. Software release at <http://cristal.inria.fr/~simonet/soft/flowcaml/>, July 2003.
- [35] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. Flexible Dynamic Information Flow Control in Haskell. In *Haskell Symposium*. ACM SIGPLAN, September 2011.
- [36] D. Stefan, A. Russo, P. Buiras, A. Levy, J. C. Mitchell, and D. Mazières. Addressing covert termination and timing channels in concurrent information flow systems. In *Proc. of the 17th ACM SIGPLAN International Conference on Functional Programming*, Sep. 2012.
- [37] D. Stefan, A. Russo, D. Mazières, and J. C. Mitchell. Disjunction category labels. In *Proceedings of the 16th Nordic Conference on Information Security Technology for Applications*, NordSec'11. Springer-Verlag, 2012.
- [38] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. Flexible dynamic information flow control in the presence of exceptions. <http://www.cse.chalmers.se/~russo/jfp15.pdf>, 2012.
- [39] D. Stefan, E. Z. Yang, P. Marchenko, A. Russo, D. Herman, B. Karp, and D. Mazières. Protecting users by confining JavaScript with COWL. In *Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX, October 2014.

$$\begin{array}{c}
\text{APP} \\
\hline
\textcolor{red}{E}[(\lambda x.t_1) t_2] \longrightarrow \textcolor{red}{E}[\{t_2 / x\} t_1] \\
\\
\text{FIX} \\
\hline
\textcolor{red}{E}[\mathbf{fix}(\lambda x.t)] \longrightarrow \textcolor{red}{E}[\{\mathbf{fix}(\lambda x.t) / x\} t] \\
\\
\text{IFTRUE} \qquad \qquad \text{IFFALSE} \\
\hline
\textcolor{red}{E}[\mathbf{if True then } t_2 \mathbf{ else } t_3] \longrightarrow \textcolor{red}{E}[t_2] \qquad \textcolor{red}{E}[\mathbf{if False then } t_2 \mathbf{ else } t_3] \longrightarrow \textcolor{red}{E}[t_3] \\
\\
\text{LABELOP} \qquad \qquad \text{RETURN} \\
\hline
\textcolor{red}{E}[l_1 \otimes l_2] \longrightarrow \textcolor{red}{E}[v] \qquad \langle \Sigma | \textcolor{blue}{E}[\mathbf{return } t] \rangle \longrightarrow \langle \Sigma | \textcolor{blue}{E}[LIO^{\text{TCB}} t] \rangle \\
\text{where } v = \llbracket l_1 \otimes l_2 \rrbracket_\ell \\
\\
\text{BIND} \\
\hline
\langle \Sigma | \textcolor{blue}{E}[(LIO^{\text{TCB}} t_1) \gg t_2] \rangle \longrightarrow \langle \Sigma | \textcolor{blue}{E}[t_2 t_1] \rangle
\end{array}$$

Fig. 19: Reduction rules for standard  $\lambda_\ell^{\text{uo}}$  terms.

- [40] P. Wadler. Monads for functional programming. In M. Broy, editor, *Marktoberdorf Summer School on Program Design Calculi*, volume 118 of *NATO ASI Series F: Computer and systems sciences*. Springer Verlag, August 1992.
- [41] E. Yang, D. Stefan, J. Mitchell, D. Mazières, P. Marchenko, and B. Karp. Toward principled browser security. In *The 14th Workshop on Hot Topics in Operating Systems (HotOS XIV)*. USENIX, 2013.
- [42] A. Yip, N. Narula, M. Krohn, and R. Morris. Privacy-preserving browser-side scripting with bflow. In *Proc. of the 4th ACM European Conference on Computer Systems*, EuroSys '09. ACM, 2009.
- [43] S. Zdancewic. PhD thesis: Programming languages for information security. Technical report, Cornell University, 2002.
- [44] N. Zeldovich, S. Boyd-Wickizer, E. Kohler, and D. Mazières. Making information flow explicit in HiStar. In *Proc. of the 7th Symp. on Operating Systems Design and Implementation*, pages 263–278, Seattle, WA, November 2006.

## A Semantics for the base calculus

The reduction rules for pure and monadic terms are given in Figure 19. We define substitution  $\{t_2 / x\} t_1$  in the usual way: homomorphic on all operators and renaming bound names to avoid captures. Our label operations  $\sqcup$ ,  $\sqcap$ , and  $\sqsubseteq$  rely on the label-specific implementation of these lattice operators, as used in the premise of rule (LABELOP); we use the meta-level partial function  $\llbracket \cdot \rrbracket_\ell$ , which maps terms to values, to precisely capture this implementation detail.

The reduction rules for pure terms are standard. For instance, in rule (IFTRUE), when the branch has a true condition, i.e.,  $\textcolor{red}{E}[\mathbf{if True then } t_2 \mathbf{ else } t_3]$ , it reduces to the then branch ( $\textcolor{red}{E}[t_2]$ ). The rest are self-explanatory and we do not discuss them any further.

Since all the IFC checks are performed by individual LIO terms, the definition for **return** and ( $\gg$ ) are trivial. The former simply reduces to a monadic value by wrapping the term with the  $LIO^{\text{TCB}}$  constructor, while the latter evaluates the left-hand term and supplies the result to the right-hand term, as usual.

## B Attack on naive flow-sensitive references

```

leakRef :: RefFS Bool → LIO Bool
leakRef href = do
  lref ← newRef L True
  tmp ← newRef L False
  toLabeled H $ do h ← readRef href
                  when h $ writeRef tmp True
  toLabeled H $ do t ← readRef tmp
                  when (¬ t) $ writeRef lref False
  readRef lref

```

Fig. 20: An attack in LIO with naive flow-sensitive reference extension without **labelOf**.

As in the attack of Figure 7, the *leakRef* of Figure 20 can be used to leak the value stored in a **H** reference *href*, while keeping the current label **L**, without using **labelOf**. Internally, the value is leaked into public reference *lref* by leveraging the fact that, based on a secret value, the label of a public reference (*tmp*) can be changed (or not). In the first **toLabeled** block, if  $h \equiv \text{True}$ , then the label of *tmp* is raised to **H** and its value is set to *True*. In the second **toLabeled** block, we read *tmp*, which may raise the current label to **H** if the secret is *True* (and thus *tmp* was upgraded). Indeed, if the secret is *True* (and thus  $t \equiv \text{True}$ ) we leave the public reference intact: *True*. However, if the secret is *False*, the *tmp* reference is not modified in the first **toLabeled** block and thus when reading it in the second **toLabeled** block, the current label remains **L**, and since  $t \equiv \text{False}$ , we write *False* into the public reference. In both cases the value stored in *lref* corresponds to that of *href*, yet leaving the current label and the label of *lref* intact (**L**).

## C Embedding Theorem

In this section we prove that the embedding from  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$  into  $\lambda_{\ell}^{\text{LIO}}$  preserves semantics.

We will use the following lemma for single  $\lambda_{\ell, \text{FS}}^{\text{LIO}}$  steps:



**Lemma 2 (Single-step embedding).** *Let  $t$  be a well-typed term in  $\lambda_{\ell, \text{FS}}^{\text{uo}}$ . Then if  $\langle \Sigma | t \rangle \longrightarrow \langle \Sigma' | t' \rangle$ , then there is a configuration  $Y$  such that  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket t \rrbracket_{\text{FI}}^{\Sigma} \rangle \longrightarrow^* Y$  and  $\langle \llbracket \Sigma' \rrbracket_{\text{FI}} | \llbracket t' \rrbracket_{\text{FI}}^{\Sigma'} \rangle \longrightarrow^* Y$ , i.e.  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket t \rrbracket_{\text{FI}}^{\Sigma} \rangle$  and  $\langle \llbracket \Sigma' \rrbracket_{\text{FI}} | \llbracket t' \rrbracket_{\text{FI}}^{\Sigma'} \rangle$  are  $\beta$ -equivalent.*

*Proof.* Case analysis on the next redex in  $t$ . Most cases show a stronger version of the lemma, i.e. that  $\langle \Sigma | t \rangle \longrightarrow \langle \Sigma' | t' \rangle$  implies  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket t \rrbracket_{\text{FI}}^{\Sigma} \rangle \longrightarrow^* \langle \llbracket \Sigma' \rrbracket_{\text{FI}} | \llbracket t' \rrbracket_{\text{FI}}^{\Sigma'} \rangle$ .

**Case E [newRef<sub>FS</sub>  $l$   $t$ ].**

We have  $\langle \Sigma | \mathbf{E} [\text{newRef}_{\text{FS}} \ l \ t] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\text{return} \ (Ref_{\text{FS}}^{\text{TCB}} \ a)] \rangle$ , where  $\Sigma' = \Sigma [\mu_{\text{FS}} \mapsto \Sigma.\mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} \ \Sigma.l_{\text{cur}} (Lb^{\text{TCB}} \ l \ t)]]$ , and we know that  $\Sigma.l_{\text{cur}} \sqsubseteq l$ .

Let  $\Sigma_1 = \llbracket \Sigma \rrbracket_{\text{FI}}$ . We argue

$$\begin{aligned} & \langle \Sigma_1 | \llbracket \mathbf{E} [\text{newRef}_{\text{FS}} \ l \ t] \rrbracket_{\text{FI}}^{\Sigma} \rangle \\ & \longrightarrow \langle \Sigma'_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{do } l_{\text{cur}} \leftarrow \text{getLabel}; \text{return} \ (\text{wrap} \ (Ref_{\text{FI}}^{\text{TCB}} \ l_{\text{cur}} \ a))] \rangle \\ & \quad (\Sigma'_1 = \Sigma_1 [\mu_{\text{FI}} \mapsto \Sigma_1.\mu_{\text{FI}} [i \mapsto Lb^{\text{TCB}} \ l \ t]]) \\ & \longrightarrow \langle \Sigma''_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{return} \ (\text{wrap} \ (Ref_{\text{FI}}^{\text{TCB}} \ l_{\text{cur}} \ a))] \rangle \\ & \quad (\Sigma''_1 = \Sigma_1 [\mu_{\text{FI}} \mapsto \Sigma_1.\mu_{\text{FI}} [i \mapsto Lb^{\text{TCB}} \ l \ t; \\ & \quad \quad \quad a \mapsto Lb^{\text{TCB}} \ l_{\text{cur}} (Ref_{\text{FI}}^{\text{TCB}} \ l \ i)]] \end{aligned}$$

We now have to check that  $\llbracket \text{return} \ (Ref_{\text{FS}}^{\text{TCB}} \ a) \rrbracket_{\text{FI}}^{\Sigma'} = \text{return} \ (\text{wrap} \ (Ref_{\text{FI}}^{\text{TCB}} \ l_{\text{cur}} \ a))$  and  $\llbracket \Sigma' \rrbracket_{\text{FI}} = \Sigma''_1$ , which follow directly from the definition of  $\llbracket \cdot \rrbracket_{\text{FI}}$  for references and states.

**Case E [readRef<sub>FS</sub>  $(Ref_{\text{FS}}^{\text{TCB}} \ a)$ ].**

We have  $\langle \Sigma | \mathbf{E} [\text{readRef}_{\text{FS}} \ (Ref_{\text{FS}}^{\text{TCB}} \ a)] \rangle \longrightarrow \langle \Sigma | \mathbf{E} [\text{unlabel} \ (Lb^{\text{TCB}} \ (l \sqcup l') \ t))] \rangle$ , where  $\Sigma.\mu_{\text{FS}}(a) = Lb^{\text{TCB}} \ l \ (Lb^{\text{TCB}} \ l' \ t)$ .

Let  $\Sigma_1 = \llbracket \Sigma \rrbracket_{\text{FI}}$ . We argue

$$\begin{aligned} & \langle \Sigma_1 | \llbracket \mathbf{E} [\text{readRef}_{\text{FS}} \ (Ref_{\text{FS}}^{\text{TCB}} \ a)] \rrbracket_{\text{FI}}^{\Sigma} \rangle \\ & \longrightarrow \langle \Sigma_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{readRef}_{\text{FI}} \ (\llbracket Ref_{\text{FS}}^{\text{TCB}} \ a \rrbracket_{\text{FI}}) \gg \text{readRef}_{\text{FI}}] \rangle \\ & \longrightarrow \langle \Sigma_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{unlabel} \ (\Sigma_1.\mu_{\text{FI}}(a)) \gg \text{readRef}_{\text{FI}}] \rangle \\ & \longrightarrow \langle \Sigma'_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{return} \ (Ref_{\text{FI}}^{\text{TCB}} \ l' \ i) \gg \text{readRef}_{\text{FI}}] \rangle \\ & \quad (\Sigma'_1 = \Sigma_1 [l_{\text{cur}} \mapsto \Sigma_1.l_{\text{cur}} \sqcup l]) \\ & \longrightarrow \langle \Sigma'_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{readRef}_{\text{FI}} \ (Ref_{\text{FI}}^{\text{TCB}} \ l' \ i)] \rangle \\ & \longrightarrow \langle \Sigma'_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{unlabel} \ (\Sigma'_1.\mu_{\text{FI}}(i))] \rangle \\ & \longrightarrow \langle \Sigma''_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{return} \ t] \rangle \\ & \quad (\Sigma''_1 = \Sigma'_1 [l_{\text{cur}} \mapsto \Sigma'_1.l_{\text{cur}} \sqcup l']) \end{aligned}$$

Now if we consider  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket \mathbf{E} [\text{unlabel} \ (Lb^{\text{TCB}} \ (l \sqcup l') \ t))] \rrbracket_{\text{FI}} \rangle$ , we have

$$\begin{aligned} & \langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket \mathbf{E} [\text{unlabel} \ (Lb^{\text{TCB}} \ (l \sqcup l') \ t))] \rrbracket_{\text{FI}} \rangle \\ & \longrightarrow \langle \Sigma_2 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\text{return} \ t] \rangle \\ & \quad (\Sigma_2 = \Sigma_2 [l_{\text{cur}} \mapsto (\llbracket \Sigma \rrbracket_{\text{FI}}).l_{\text{cur}} \sqcup l \sqcup l']) \end{aligned}$$

Note that  $\Sigma''_1.l_{\text{cur}} = (\llbracket \Sigma \rrbracket_{\text{FI}}).l_{\text{cur}} \sqcup l \sqcup l' = \Sigma_2.l_{\text{cur}}$ .

**Case E [writeRef<sub>FS</sub>  $(Ref_{\text{FS}}^{\text{TCB}} \ a) \ t$ ].**

We have  $\langle \Sigma | \mathbf{E} [\mathbf{writeRef}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) t] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\mathbf{return} ()] \rangle$ , where  $\Sigma.\mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l' v)$ ,  $\Sigma' = \Sigma [\mu_{\text{FS}} \mapsto \Sigma.\mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} l (Lb^{\text{TCB}} l' v)]]$ , and we know that  $\Sigma.l_{\text{cur}} \sqsubseteq l \sqcup l'$ .

Let  $\Sigma_1 = \llbracket \Sigma \rrbracket_{\text{FI}}$ . Then there exists a function  $\mu$  such that:

$$\begin{aligned} & \langle \Sigma_1 | \llbracket \mathbf{E} [\mathbf{upgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) l'] \rrbracket_{\text{FI}}^{\Sigma} \rangle \\ & \longrightarrow \langle \Sigma_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\mathbf{toLabeled} (\Sigma_1.l_{\text{cur}} \sqcup l) (\mu \Sigma_1.l_{\text{cur}})] \rangle \end{aligned}$$

We now step through the evaluation of  $\langle \Sigma_1 | \mu \Sigma_1.l_{\text{cur}} \rangle$ , as follows:

$$\begin{aligned} & \langle \Sigma_1 | \mu \Sigma_1.l_{\text{cur}} \rangle \\ & \longrightarrow \langle \Sigma_1 | \mathbf{do} i \leftarrow \mathbf{readRef}_{\llbracket Ref_{\text{FS}}^{\text{TCB}} a \rrbracket_{\text{FI}}}; \mathbf{writeRef}_{\text{FI}} i t \rangle \\ & \longrightarrow \langle \Sigma'_1 | \mathbf{writeRef}_{\text{FI}} i t \rangle \\ & \quad (\Sigma'_1 = \Sigma_1 [l_{\text{cur}} \mapsto \Sigma_1.l_{\text{cur}} \sqcup l]) \\ & \longrightarrow \langle \Sigma''_1 | \mathbf{return} () \rangle \\ & \quad (\Sigma''_1 = \Sigma'_1 [\mu_{\text{FI}} \mapsto \Sigma'_1.\mu_{\text{FI}} [i \mapsto Lb^{\text{TCB}} l' t]]) \end{aligned}$$

Finally, this allows us to conclude (from the rule for **toLabeled**), that

$$\begin{aligned} & \langle \Sigma_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\mathbf{toLabeled} (\Sigma_1.l_{\text{cur}} \sqcup l) (\mu \Sigma_1.l_{\text{cur}})] \rangle \\ & \longrightarrow \langle \Sigma_2 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\mathbf{return} ()] \rangle \end{aligned}$$

where  $\Sigma_2 = (\Sigma_1.l_{\text{cur}}, \Sigma''_1.\mu_{\text{FI}})$ . Now we can check that  $\llbracket \Sigma' \rrbracket_{\text{FI}} = \Sigma_2$  from the definition of  $\llbracket \cdot \rrbracket_{\text{FI}}$  for states.

**Case E**  $[\mathbf{upgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) l']$ .

We have  $\langle \Sigma | \mathbf{E} [\mathbf{upgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) l'] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\mathbf{return} ()] \rangle$ , where  $\Sigma.\mu_{\text{FS}}(a) = Lb^{\text{TCB}} l (Lb^{\text{TCB}} l'' v)$ ,  $\Sigma' = \Sigma [\mu_{\text{FS}} \mapsto \Sigma.\mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} l (Lb^{\text{TCB}} l'' v)]]$ , and we know that  $\Sigma.l_{\text{cur}} \sqsubseteq l$ .

Let  $\Sigma_1 = \llbracket \Sigma \rrbracket_{\text{FI}}$ . Then there exists a function  $\mu$  such that:

$$\begin{aligned} & \langle \Sigma_1 | \llbracket \mathbf{E} [\mathbf{upgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) l'] \rrbracket_{\text{FI}}^{\Sigma} \rangle \\ & \longrightarrow \langle \Sigma_1 | (\llbracket \mathbf{E} \rrbracket_{\text{FI}}^{\Sigma}) [\mathbf{toLabeled} (\Sigma_1.l_{\text{cur}} \sqcup l) (\mu \Sigma_1.l_{\text{cur}})] \rangle \end{aligned}$$

We now step through the evaluation of  $\langle \Sigma_1 | \mu \Sigma_1.l_{\text{cur}} \rangle$ , as follows:

$$\begin{aligned} & \langle \Sigma_1 | \mu \Sigma_1.l_{\text{cur}} \rangle \\ & \longrightarrow \langle \Sigma_1 | \mathbf{do} i \leftarrow \mathbf{readRef}_{\llbracket Ref_{\text{FS}}^{\text{TCB}} a \rrbracket_{\text{FI}}}; lc \leftarrow \mathbf{getLabel}; \dots \rangle \\ & \longrightarrow \langle \Sigma'_1 | \mathbf{do} lc \leftarrow \mathbf{getLabel}; n \leftarrow \mathbf{newRef}_{\text{FI}} (lc \sqcup (l' \sqcup l)) \perp; \dots \rangle \\ & \quad (\Sigma'_1 = \Sigma_1 [l_{\text{cur}} \mapsto \Sigma_1.l_{\text{cur}} \sqcup l]) \\ & \longrightarrow \langle \Sigma'_1 | \mathbf{do} n \leftarrow \mathbf{newRef}_{\text{FI}} (lc \sqcup (l' \sqcup l)) \perp; \mathbf{copyRef} i n; \dots \rangle \\ & \longrightarrow \langle \Sigma''_1 | \mathbf{do} \mathbf{copyRef} i n; \mathbf{writeRef}_{\text{FI}} (\llbracket Ref_{\text{FS}}^{\text{TCB}} a \rrbracket_{\text{FI}}) n \rangle \\ & \quad (\Sigma''_1 = \Sigma'_1 [\mu_{\text{FI}} \mapsto \Sigma'_1.\mu_{\text{FI}} [n \mapsto Lb^{\text{TCB}} (lc \sqcup (l' \sqcup l)) \perp]]) \\ & \longrightarrow \langle \Sigma'''_1 | \mathbf{writeRef}_{\text{FI}} (\llbracket Ref_{\text{FS}}^{\text{TCB}} a \rrbracket_{\text{FI}}) n \rangle \\ & \quad (\Sigma'''_1 = \Sigma''_1 [\mu_{\text{FI}} \mapsto \Sigma''_1.\mu_{\text{FI}} [n \mapsto Lb^{\text{TCB}} (lc \sqcup (l' \sqcup l)) v]]) \\ & \longrightarrow \langle \Sigma'''_1 | \mathbf{return} () \rangle \\ & \quad (\Sigma'''_1 = \Sigma''_1 [\mu_{\text{FI}} \mapsto \Sigma''_1.\mu_{\text{FI}} [a \mapsto Lb^{\text{TCB}} l (Ref_{\text{FI}}^{\text{TCB}} (lc \sqcup l' \sqcup l) n)]] \end{aligned}$$

Finally, this allows us to conclude (from the rule for **toLabeled**), that

$$\begin{aligned} & \langle \Sigma_1 | ([\mathbf{E}]_{\text{FI}}^\Sigma) [\mathbf{toLabeled} \, l' \, (\mu \, \Sigma_1.l_{\text{cur}})] \rangle \\ & \longrightarrow \langle \Sigma_2 | ([\mathbf{E}]_{\text{FI}}^\Sigma) [\mathbf{return} \, ()] \rangle \end{aligned}$$

where  $\Sigma_2 = (\Sigma_1.l_{\text{cur}}, \Sigma_1'''.\mu_{\text{FI}})$ . Now we can check that  $\llbracket \Sigma' \rrbracket_{\text{FI}} = \Sigma_2$  from the definition of  $\llbracket \cdot \rrbracket_{\text{FI}}$  for states.

**Case E**  $[\mathbf{downgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) \, l']$ .

We have  $\langle \Sigma | \mathbf{E} [\mathbf{downgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) \, l'] \rangle \longrightarrow \langle \Sigma' | \mathbf{E} [\mathbf{return} \, ()] \rangle$ , where  $\Sigma.\mu_{\text{FS}}(a) = Lb^{\text{TCB}} \, l \, (Lb^{\text{TCB}} \, l'' \, v)$ ,  $\Sigma' = \Sigma [\mu_{\text{FS}} \mapsto \Sigma.\mu_{\text{FS}} [a \mapsto Lb^{\text{TCB}} \, l \, (Lb^{\text{TCB}} (l \sqcup l'' \sqcap l') \perp)]]$ , and we know that  $\Sigma.l_{\text{cur}} \sqsubseteq l$ .

Let  $\Sigma_1 = \llbracket \Sigma \rrbracket_{\text{FI}}$ . Then there exists a function  $\mu$  such that:

$$\begin{aligned} & \langle \Sigma_1 | [\mathbf{E} [\mathbf{downgrade}_{\text{FS}} (Ref_{\text{FS}}^{\text{TCB}} a) \, l']]_{\text{FI}}^\Sigma \rangle \\ & \longrightarrow \langle \Sigma_1 | ([\mathbf{E}]_{\text{FI}}^\Sigma) [\mathbf{toLabeled} \, (\Sigma_1.l_{\text{cur}} \sqcup l) \, (\mu \, \Sigma_1.l_{\text{cur}})] \rangle \end{aligned}$$

We now step through the evaluation of  $\langle \Sigma_1 | \mu \, \Sigma_1.l_{\text{cur}} \rangle$ , as follows:

$$\begin{aligned} & \langle \Sigma_1 | \mu \, \Sigma_1.l_{\text{cur}} \rangle \\ & \longrightarrow \langle \Sigma_1 | \mathbf{do} \, i \leftarrow \mathbf{readRef}_{\llbracket Ref_{\text{FS}}^{\text{TCB}} a \rrbracket_{\text{FI}}} ; lc \leftarrow \mathbf{getlabel}; \dots \rangle \\ & \longrightarrow \langle \Sigma'_1 | \mathbf{do} \, lc \leftarrow \mathbf{getLabel}; n \leftarrow \mathbf{newRef}_{\text{FI}} (lc \sqcup (l' \sqcap l)) \perp; \dots \rangle \\ & \quad (\Sigma'_1 = \Sigma_1 [l_{\text{cur}} \mapsto \Sigma_1.l_{\text{cur}} \sqcup l]) \\ & \longrightarrow \langle \Sigma'_1 | \mathbf{do} \, n \leftarrow \mathbf{newRef}_{\text{FI}} (lc \sqcup (l' \sqcap l)) \perp; \mathbf{writeRef}_{\text{FI}} (\llbracket Ref_{\text{FS}}^{\text{TCB}} a \rrbracket_{\text{FI}}) \, n \rangle \\ & \longrightarrow \langle \Sigma''_1 | \mathbf{writeRef}_{\text{FI}} (\llbracket Ref_{\text{FS}}^{\text{TCB}} a \rrbracket_{\text{FI}}) \, n \rangle \\ & \quad (\Sigma''_1 = \Sigma'_1 [\mu_{\text{FI}} \mapsto \Sigma'_1.\mu_{\text{FI}} [n \mapsto Lb^{\text{TCB}} (lc \sqcup (l' \sqcap l)) \perp]]) \\ & \longrightarrow \langle \Sigma'''_1 | \mathbf{return} \, () \rangle \end{aligned}$$

Finally, this allows us to conclude (from the rule for **toLabeled**), that

$$\begin{aligned} & \langle \Sigma_1 | ([\mathbf{E}]_{\text{FI}}^\Sigma) [\mathbf{toLabeled} \, l' \, (\mu \, \Sigma_1.l_{\text{cur}})] \rangle \\ & \longrightarrow \langle \Sigma_2 | ([\mathbf{E}]_{\text{FI}}^\Sigma) [\mathbf{return} \, ()] \rangle \end{aligned}$$

where  $\Sigma_2 = (\Sigma_1.l_{\text{cur}}, \Sigma_1'''.\mu_{\text{FI}})$ . Now we can check that  $\llbracket \Sigma' \rrbracket_{\text{FI}} = \Sigma_2$  from the definition of  $\llbracket \cdot \rrbracket_{\text{FI}}$  for states.

Now we can state the main theorem of this section.

**Theorem.** [Embedding  $\lambda_{\ell, \text{FS}}^{\text{uo}}$  in  $\lambda_\ell^{\text{uo}}$ ] Let  $t$  be a well-typed term in  $\lambda_{\ell, \text{FS}}^{\text{uo}}$ . Then if  $\langle \Sigma | t \rangle \longrightarrow^* \langle \Sigma' | v \rangle$ , we have  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket t \rrbracket_{\text{FI}}^\Sigma \rangle \longrightarrow^* \langle \llbracket \Sigma' \rrbracket_{\text{FI}} | \llbracket v \rrbracket_{\text{FI}}^{\Sigma'} \rangle$ , and if  $\langle \Sigma | t \rangle \longrightarrow^* \langle \Sigma' | \uparrow \rangle$ , then  $\langle \llbracket \Sigma \rrbracket_{\text{FI}} | \llbracket t \rrbracket_{\text{FI}}^\Sigma \rangle \longrightarrow^* \langle \llbracket \Sigma' \rrbracket_{\text{FI}} | \uparrow \rangle$ .

*Proof.* By induction on the number of steps in  $\langle \Sigma | t \rangle \longrightarrow^* \langle \Sigma' | v \rangle$ , using Lemma 2 and uniqueness of normal forms in  $\lambda_\ell^{\text{uo}}$ .